Canadian Journal of Philosophy (2026), 56: e24, 1–17
doi:10.1017/can.2026.10073

SYMPOSIUM ARTICLE

Nettles in the Nexus: On Wallace on Relational Morality

Gopal Sreenivasan

Department of Philosophy, Duke University, Durham, NC, USA
Email: gopal.sreenivasan@duke.edu

Abstract
Wallace aims to interpret all of morality in relational terms. The paradigmatic features of his relational interpretation derive from the case of promising. I agree both that morality has a rich and interesting relational core and that promises illustrate it well. Against Wallace, however, I argue that other significant elements of morality resist assimilation to his relational interpretation. In some cases, this is because their relational character differs notably from that of promising. In other cases, important elements of morality lack a relational character altogether. Wallace's interpretation therefore stands to be improved by sharply reducing the scope of its ambition.

Keywords: relational morality; claims; rights; directed duties; wrongings; moderate deontology

In his magisterial work, *The Moral Nexus*, Jay Wallace proposes to interpret all of morality in relational terms (Wallace, 2019).[1] As Wallace explains, his relational interpretation of morality has three core features, all of which are present in ordinary cases of promising (pp. 5–11).[2] To give ourselves an example, suppose I promise my neighbor, Mia, that I will paint her fence next week. In that case, the following propositions are licensed uncontroversially:

(i). I have a duty *to Mia* to paint her fence next week;
(ii). Mia has a claim-right *against me* that I paint her fence next week;
(iii). if I violate my duty to paint Mia's fence, then I *wrong her*.

Each of these propositions has a relational dimension, which has been flagged in italics. For instance, the relational dimension of the first proposition resides in the fact that my promissory obligation is itself a relational obligation. Alternatively, in the language of duties, it is a directed duty. That is to say, instead of being a duty that the duty-bearer (in this case, me) has all by himself, a directed duty always holds between two parties, whom it binds together, the duty-bearer and the person to whom the duty is owed (here, my neighbor).[3]

[1]It may be more accurate to say that he proposes to interpret "almost all" of morality in relational terms. But for introductory purposes, a crisp and simple statement of Wallace's ambition works better. We shall attend to some of the relevant complications in due course (Section 4).

[2]Bare page citations all refer to Wallace (2019).

[3]I shall return to the relational dimension of the other two propositions in short order.

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Canadian Journal of Philosophy, Inc. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives licence (http://creativecommons.org/licenses/by-nc-nd/4.0), which permits non-commercial re-use, distribution, and reproduction in any medium, provided that no alterations are made and the original article is properly cited. The written permission of Cambridge University Press or the rights holder(s) must be obtained prior to any commercial use and/or adaptation of the article.

One useful way to understand the arc of Wallace's central argument is that he seeks to extend the nexus of relational features described by propositions (i)–(iii) outward from the basic case of promising to the rest of morality (pp. 15–16). I accept both that Wallace's relational nexus fits the case of promising very well and that it extends smoothly to considerably more of morality than promising alone. However, I shall argue that there are also significant elements of morality to which the relational nexus does *not* extend, so that Wallace's ambition overshoots the mark in certain important respects. Fortunately, moreover, this conclusion can be established without having to ascertain exactly where the outer boundary of that ambition lies. This argument, which concerns the *scope* of the relational nexus, will constitute one axis of my critique.

A second axis of my critique will concern two *extraneous* features Wallace builds into his account of the relational nexus. According to Wallace, directed duties are both unanalyzable (p. 16) and absolute (pp. 170–72). By "unanalyzable," I mean that a directed duty cannot be factored into components (e.g., in an attempt to shed light on what it means for a duty to have a direction).[4] Instead, for Wallace, the fact that promissory duties have a direction is simply a theoretical primitive (p. 159). By "absolute," I mean, following Thomson (1990), that my duty (e.g., to paint Mia's fence next week) *entails* that, all things considered (ATC), I ought to paint Mia's fence next week. While Wallace imputes both of these features to promissory duties, they are extraneous to the relational nexus per se, insofar as neither is necessary for promissory duties to have a direction. Hence, neither feature is necessary to account for the *relational* character of promissory duties, from which paradigm the nexus acquires its eponymous character.

I reject both of these extraneous features. However, this is not because they are extraneous, but rather because they yield a false picture of most moral duties. On this second axis of my critique, the argument sometimes proceeds directly and other times indirectly. My direct argument against Wallace targets his thesis that directed duties are absolute. I shall argue that some duties are not absolute. On the assumption that all duties are directed duties,[5] it follows that some directed duties are not absolute.

Like some of the other controversies we shall encounter in connection with Wallace's extraneous features, the question of whether directed duties are absolute is the subject of a long-standing debate. Notably, for example, Judith Thomson (1990) argues that directed duties are not absolute. I myself have always found her argument thoroughly convincing, but Wallace clearly has a different opinion.[6] However, in this paper, I shall strive to avoid re-litigating familiar arguments (or even reviewing them). In their place—or perhaps, by way of supplement—I advance a novel argument against Wallace's thesis.

Finally, the two axes of my critique also interact to significant effect. For example, their interaction drives my indirect argument against the proposition that directed duties are unanalyzable. As we shall see, that proposition will prove to be an obstacle to assimilating sundry instances of moral wronging to core features of Wallace's paradigm case. Since assimilation to promising is the mechanism by which the perimeter of his relational nexus gets extended, anything that prevents the wrongs in question from being assimilated to promissory wrongs strands them on the far side of that perimeter. Wallace is accordingly vulnerable to a dilemma. Either directed duties are analyzable after all or various kinds of wronging must remain outside the relational nexus.

At its core, then, Wallace's portrait of morality as a relational nexus is valuable and illuminating in equal measure. Past a certain point, however, defects in this portrait proliferate rapidly, as I shall now endeavor to demonstrate. These are the nettles in the nexus to which my title refers.

[4]In particular, Wallace rejects the idea that a directed duty can be factored into an undirected duty and something additional. He explicitly disavows the attempts prosecuted along these lines by the will theory and the interest theory of rights (pp. 157–159).

[5]Wallace is committed to this assumption, even though he does not argue for it explicitly. For any undirected moral duty—that is, any duty that is not owed to anyone—is automatically ejected from the relational nexus, and hence escapes Wallace's ambition. In Section 2, I shall argue that this assumption is false.

[6]For a tidy reprise of Wallace's reasons for resisting Thomson's position, see Wallace (2026).

1. Varieties of Moral Wronging

Let us start with the phenomenon of moral wronging. Wronging is an inherently relational notion. To that extent, it should be grist to Wallace's mill. What makes wronging relational is the fact that it always involves a second party, someone else who is wronged and whom the agent wrongs. By contrast, one can act wrongly all by oneself. In our opening promising example, this contrast between wronging and acting wrongly is borne out by proposition (iii), on which Mia would be wronged by my failure to paint her fence next week. More generally, in promising, the promisee will be wronged if the promisor fails to perform on the promise.

A lot of recent philosophical attention to wronging focuses on the apparent diversity in ordinary cases of wronging. One of the central issues in this literature concerns whether someone can be wronged by an action even if they have no claim-right against the agent not to act in this way. For example, Nicolas Cornell (2015) argues that wrongs can indeed come apart from rights, in support of which he adduces a trove of cases featuring what can be called third-party wrongs.[7] In such cases, the agent violates a second party's claim-right (thereby wronging them), and by this same action wrongs a third party as well, although this third party lacks any claim-right against the agent. Later, Cornell (2025) further diversifies his repertoire of counterexamples to include cases where no one's claim-right is violated at all. I shall eventually concentrate on one of his later cases, since it suits our purposes better. At that point, I shall color in some details of the case. For the moment, however, let us make do with the bare structure of a third-party wrong.

That structure is a useful place to begin, since reflecting on it serves to clarify three crucial preliminary points. First of all, at least on the surface, third-party wrongs differ significantly from promissory wrongs. As proposition (ii) under our opening example illustrates, promisees have a claim-right against the promisor to the performance of their promise. Unlike third-party wrongs, then, promissory wrongs do *not* come apart from claim-rights. Third-party wrongs therefore seem not to exemplify the relational nexus in its full glory, despite their relational character.

Second, if we follow convention and adopt Wesley Hohfeld's (1919) conception of a claim-right, it turns out that propositions (i) and (ii) are simply two equivalent ways of reporting the same fact. According to what we can call the Hohfeldian Correlativity Axiom,

(HCA) Y has a claim-right against X that X φ if and only if X owes Y a duty that X φ.

On (HCA), the relational dimension of proposition (ii) is actually the same as that of proposition (i), only it holds in reverse: the same two parties are related by the same bond, which shows up on one end as a directed duty (promisor) and on the other end as a claim-right (promisee).

Described in these terms, the scope of the relational nexus straightforwardly extends at least as far as the realm of claim-rights generally. Wallace may have centered his nexus on promises, but its core features all accompany *any and every* moral claim-right, and not simply promissory claim-rights.

This brings us to the third crucial point. For his own part, Wallace does not employ Hohfeld's conception of a claim-right. In its place, he substitutes something else, which he calls a "claim."[8] Wallace's claims have the same structure as Hohfeldian claim-rights, but they have a wider extension. We can render the structure Wallace has in mind simply enough:

(W) Y has a *claim* against X that X φ if and only if X owes Y a duty that X φ.

While every Hohfeldian (moral) claim-right counts as one of Wallace's claims,[9] Wallace importantly affirms various kinds of claim where there is no Hohfeldian claim-right.

---

[7] His leading example is the famous case of *Palsgraf v Long Island Railroad Company* (1928).

[8] As it happens, the wording in Hohfeld's original text is also "claim." He even makes a point, somewhat curiously, of applauding the fact that it is monosyllabic. But the contemporary convention is to write Hohfeld's "claim" out as "claim-right." To forestall confusion, I shall reserve the term "claim" to refer to Wallace's sui generis notion.

[9] I am disregarding some complications here, which arise from the question of whether Hohfeldian claim-rights should be understood to have all of the extraneous features that claims have.

Some of the additional width in the extension of "claim" is intentionally geared to disarming putative examples of wrongs and [claims] coming apart. Third-party wrongs are precisely a case in point (pp. 196–200). Wallace finds many of Cornell's particular examples unconvincing. Nevertheless, he accepts that there are some genuine third-party wrongs. Still, in his view, none counts as a case in which wrongs come apart from claims. For Wallace suggests that any wronged third party can always be credited with a "secondary" claim against the agent. A secondary claim is distinguished from a primary claim (as well as from a Hohfeldian claim-right) by the fact that its holder can only be identified ex post, that is, subsequent to the violation of someone else's primary claim. Thus, Wallace regards third-party wrongs as cases where the wrong does come apart from a primary claim, but not from a secondary claim. However, since he understands "claim" to mean "primary or secondary claim," it remains the case that third-party wrongs do not come apart from claims.

We can now clarify that Wallace actually interprets the correlative of a promissory duty as a claim, too, rather than a Hohfeldian claim-right. Officially speaking, then, the second of the three core features that constitute Wallace's relational nexus—proposition (ii) under our opening example, which recognizes Mia's claim-right—should be expressed in the language of "claims." So construed, third-party wrongs can be *fully* assimilated to promissory wrongs, despite the initial appearances, which places them fully inside the relational nexus.

There are a number of difficulties with the device of the secondary claim, as well as with its corollary, the disjunctive understanding of "claims." But instead of trying to articulate them, I turn immediately to one of Cornell's later examples. It should not have escaped notice that Wallace's secondary claim reply seeks to exploit a concession inherent in the structure of a third-party wrong, namely, that it always involves the violation of someone's primary claim (albeit someone else's). The later examples I have in mind do not offer that concession.

Cornell's (2025) best example involves a case of personal betrayal, which he eruditely borrows from *War and Peace.* Natasha is secretly betrothed to Prince Andrey. During a year's delay in their wedding, instituted at Andrey's request (to placate his father), Natasha succumbs to the attentions of a scoundrel and winds up foolishly breaking her engagement to Andrey, who is devastated. Intuitively, it is very plausible not only that Natasha acts wrongly in betraying Andrey's trust, but furthermore that she wrongs him. Nevertheless, as Tolstoy makes clear, Andrey has no *right* against Natasha that she not break their engagement (in Wallace's terms, Andrey has no such claim). It is part of the story that he has explicitly given her perfect freedom to consider her options during the year's delay. (More academically, Andrey has waived his relevant rights against Natasha.) For Cornell, this is another case where wrongs and rights come apart.

I said that it is intuitively very plausible that Natasha wrongs Andrey, that her betrayal of his trust is a wronging, in the relational sense. While that is true, it is not a point on which I need to rely. For it is at least equally plausible that her betrayal was (simply) wrong.[10] That is one of the nice things about this example. However, Wallace cannot allow that Natasha's betrayal is simply wrong, since that would instantly eject her case from his nexus. (Indeed, more generally, he cannot allow that betrayals are simply wrong.) This brings out, and underlines, the considerable strength of Wallace's ambition—"strength" in the logical sense, which is not necessarily an advantage. There are no simple wrongs in the relational nexus, only wrongings (and so, in his view, none in morality either). Hence, to accommodate the intuitive wrongfulness of Natasha's betrayal, Wallace must treat it as a wronging. Dialectically, I can have that detail for free.

Another nice feature of Cornell's Tolstoy example is that we can give a clear and convincing account of the relational character of the wronging it contains. All wrongings involve both a

[10] Arguably, betrayals are actually wrong by definition (at least, pro tanto). Betrayal is a thick concept, like cruelty or murder. From this perspective, the question about Natasha is more perspicuously framed as whether succumbing to Kuragin's attentions was, precisely, a *betrayal.* However, since Natasha cheats with Kuragin *before* she breaks her engagement to Andrey, it is clear that she has acted wrongly. Equivalently, since Andrey is trusting her to keep faith with the monogamous character of an engagement (at least, as long as they remain engaged), she does indeed betray him.

wrongful act and some significant relation between the wrongdoer and a second person, whom they wrong. By an account of a wronging's "relational character," I mean an account of the significant relation binding its wrongdoer and wronged party together. To account for this relation in Natasha and Andrey's case, we can adapt a persuasive analysis developed by Adrienne Martin (2021). A little adaptation is needed because what Martin aims to explain is the relation underlying some directed duties, rather than the relation underlying some wrongings. In other words, she aims to explain what "directs" some duties *to* another person, rather than what "directs" or "attaches" the wrongfulness of some wrong acts *to* another person. Nevertheless, I suggest that the basic account she offers can supply both kinds of explanation.

Martin's (2021) account is restricted to cases of significant personal relationships. Her fundamental insight is that some of these relationships are governed by moral norms that play a dual role (cf. note 14). On the one hand, the operative norm is partly constitutive of the relationship in question. Martin's leading example is the parent–child relationship. On the other hand, the same norm imposes certain duties on one of the parties in the relationship.[11] In these cases, Martin maintains, the norm does not merely impose a duty on the duty-bearer, but also *directs* this duty *to the party* with whom the duty-bearer shares the relationship (i.e., the very relationship that norm partially constitutes at the same time). This is entirely plausible.

How much Martin's framework needs to be adapted to account for the relational character of Natasha's wronging depends on whether the wrongfulness of an action presupposes that some *duty* has been contravened.[12] If a duty *is* presupposed, no adaptation is required. For then the wrongfulness of Natasha's betrayal already implies some duty on her part not to betray Andrey. The only relevant question remaining is whether the norm that imposes this duty is also partially constitutive of Natasha's significant prior relationship with Andrey. But even to ask this question is to apply Martin's framework "as is" to the case at hand. Moreover, the answer is very plausibly "yes." A duty of monogamy is partially constitutive of the character of an engagement (at least, a conventional one). In that case, the relation underlying the duty Natasha owes to Andrey and the relation underlying her wronging of him are one and the same. Under the circumstances, our explanandum and Martin's are not distinct after all.

It may be, however, that Natasha's wronging does not have to be mediated by any duty on her part. On this alternative, the wrongfulness of her non-monogamy can follow directly from its flouting some underlying norm governing her significant prior relationship with Andrey. As before, it remains to ask whether the very norm that she flouts is also partially constitutive of their relationship. However, there is no reason to answer differently just because a mediating role for "duties" in wronging has been discarded. As long as the norm of monogamy plays this same dual role, it is still very plausible to understand the wrongfulness of an act that flouts a relationship-constituting norm as being *attached to the other party* who shares the relationship constituted by that norm. In Natasha's case, the wrongfulness of her non-monogamy is thereby attached to her fiancé, Andrey, making it the case that she wrongs him.

Wallace considers and rejects a more superficial version of this account (pp. 17; 192–5), which merely ties the direction of a directed duty to an antecedent relationship. As we shall see, he is both right and wrong to reject it, though for reasons having nothing to do with the superficiality of his nominal target. Let me first explain why Wallace is right to reject the superficial account he considers, and so would be right to reject Martin's account too. In a word, her account has the wrong scope to account for all directed duties. More specifically, it even has the wrong scope to account for all promissory duties or all duties correlative to some Hohfeldian claim-right or other.

[11]Depending on the relationship, the duties in question may or may not be imposed symmetrically on both parties. But even when a duty is imposed symmetrically, we can always elect to focus on one of the duty-bearers in particular. By contrast to the parent–child case (with which Martin is concerned), symmetry across duties is much more plausible with an engagement.

[12]Since the same result can be achieved either way, we can leave this question open.

Holders of Hohfeldian claim-rights can be complete and utter strangers to the relevant duty bearer, as can some promisees. Notably, some Hohfeldian claim-rights are held "against the world," which guarantees that many of their correlative duty bearers have no antecedent relationship with the right holder.[13] Evidently, however, Martin's account does not apply to duties owed to strangers, since it presupposes a significant prior relationship between the parties. Hence, it cannot account for all the directed duties in the core of the relational nexus, let alone for the full measure of Wallace's ambition.

Martin's account clearly does not suit Wallace's purposes. It does not follow, however, that it does not suit her own purposes. Still less does it follow that Martin's account is not perfectly in order when explicitly restricted to the context of significant personal relationships.[14] In particular, the required limitation on its scope does not prevent it from explaining the relational character of Natasha's wronging, whose case fits easily within that limited scope.

On the face of things, Natasha's wronging escapes the reach of Wallace's ambition for the same reason that third-party wrongs appeared to escape it, namely, not being fully assimilable to promissory wrongs. Specifically, across this board, the wronged party has no claim-right against the wrongdoer, unlike wronged promisees. What is more, in Natasha's case, there is no toehold for Wallace's device of the secondary claim either. All the same, there is actually a more robust and important respect in which Natasha's wronging, and other wrongings relevantly like it, are dissimilar to promissory wrongings, and therefore excluded from the relational nexus.

To wit, the relational character of Natasha's wronging differs from that of promissory wrongings. As we have seen, the relational character of her wronging arises in the context of a significant personal relationship between the two agents who wind up as wrongdoer and wronged. For ease of reference, let us call this kind of wronging a "relationship-based" wronging. By contrast, the relational character of promissory wrongs arises from something else. This much is already fixed by the fact that this second wronging relation has a wider scope, insofar as its poles can be occupied by complete strangers.

I shall not try to account for the nature of this "something else." We can hold a place for an account of it open by observing that the relation underlying promissory wrongs precisely mirrors the relation between Hohfeldian claim-right holders and the bearers of their correlative duties, as expressed by (HCA). But whether or not that place can be filled adequately,[15] the foregoing difference in scope is enough to distinguish the relation expressed by (HCA) from the relation underlying relationship-based wrongings. To advertise this critical aspect of their relational character, we can refer to promissory wrongings as "right-based" wrongings.[16]

Although relationship-based wrongings and right-based wrongings are both relational moral phenomena—and, indeed, both kinds of wronging—they involve different underlying relations between the wrongdoer and the wronged party. Since *right-based* wrongings constitute the paradigm in terms of which Wallace defines his relational nexus, relationship-based wrongings do not belong in it. Accordingly, relationship-based wrongings are the first exhibit in my catalog of elements of morality to which Wallace's relational nexus does not extend.

---

[13]Simplifying slightly, and as the name suggests, claim-rights against the world are held against everyone else in the world. Paradigmatic examples are property rights and rights to bodily integrity. For some discussion, see Sreenivasan (2024).

[14]Ultimately, Martin's (2021) purpose is to supply counterexamples to (HCA). In this paper, I am only concerned with the ability of her account to explain the relational character of wrongings in the context of significant personal relationships. (I should perhaps also record that my discussion has abstracted from a second part of her account, which concerns reactive attitudes.)

[15]Which account, if any, of the relation expressed by (HCA) is the correct account is the question at the heart of the debate between the will theory and the interest theory of rights. For an overview of this debate, see Sreenivasan (2005), which also offers to resolve it in favor of a third option. As we have noted, Wallace eschews this debate. However, no solution to it would serve his purposes anyhow, since (W) has a wider scope in turn than (HCA).

[16]Calling them "right-based" does not prejudice Wallace's revisionary interpretation of the correlatives of promissory duties, since he understands Hohfeldian claim-rights as a "sub-class" of his claims (p. 9).

There are two responses available to Wallace here, representing diametrically opposed reactions to the apparent diversity in wronging relations on which my first exhibit rests. To begin with, he might embrace this diversity. That is, Wallace might redefine his model of the relational nexus so that any kind of relational character suffices for inclusion. Instead of privileging one specific kind of underlying moral relation, he might allow any kind of moral relation to count. The trouble with this first response is that Wallace has additional goals for his portrait of the relational nexus, such as bringing out the "underlying unity" of interpersonal morality (p. 4) and supplying a satisfying explanation of the basic phenomenon of moral obligation (p. 12 and ch. 2). Unfortunately, no nexus can accommodate relational pluralism and achieve these other goals at the same time. Nothing can be both a disjoined aggregate of heterogeneous moral relations (giving rise to different kinds of wrongings and directed duties) *and* the basis of a unifying explanation of interpersonal morality or moral obligation.[17] I therefore assume that Wallace has no interest in this first response.

Alternatively, Wallace might seek to eliminate the apparent diversity instead. For all I have said, there may well be some more penetrating or enlightened perspective from which a deeper unity in the superficially disparate relational moral phenomena is visible. That remains to be established. However, articulating this unifying perspective is incompatible with treating the directedness of a directed duty as a theoretical primitive and likewise with the relational character of claims or wrongings. Treating "claim" as a primitive is at best a placeholder for a deeper, truly unifying account. To fill that place—actually to deliver the unifying account, and thereby eliminate the apparent diversity—one must be prepared to theorize the nature of the relevant moral relation(s).

We have now come upon the dilemma for Wallace I prefigured in my introduction. Recall that one of the extraneous features he ascribes to directed duties is that they are unanalyzable. If directed duties really are unanalyzable, then the nature of the relation between a given duty-bearer and the terminus of their duty cannot be theorized. In that case, we are pre-emptively deprived of the unifying account that might eliminate the surface diversity in relational moral phenomena, including my showcase difference between relationship-based wrongings and right-based wrongings. Hence, either directed duties are analyzable or relationship-based wrongings must remain excluded from Wallace's relational nexus.

2. The ATC Ought has No Direction

The first axis of my critique targets the scope of Wallace's relational nexus. I contend that various significant elements of morality do not belong in his nexus, which cramps his ambition to expand its domain over all of morality. Exhibit one along this axis was relationship-based wrongings.

A natural place to look for further exhibits would be different precincts of morality normally regarded as lying beyond the realm of rights. On a commonsense view, there are many good examples, including what might be called the realm of virtue and various kinds of "public" duties, such as a duty to supply this or that public good and duties to steward all manner of valuable resources (notably, the earth itself). Dialectically, however, these examples are somewhat messy.

On the one hand, it is often hard to get beyond a brute clash of intuitions. For example, the last line of Thomson's book neatly encapsulates the commonsense position, "the realm of rights does not exhaust all of morality" (1990, p. 373). While Wallace opens his book by tipping his hat to her position, and that of other like-minded luminaries, he then simply rejects it (p. 2). On the other hand, the issue of which particular examples may still be worth digging in to contest is also entangled with the tricky question of where exactly the outer boundaries of Wallace's ambition lie.

[17]If that is right, a difficult question arises for Wallace concerning whether his sui generis notion of a "claim" is actually sufficiently different from a disjoined aggregate to discharge his explanatory ambitions. We have already seen that he explicitly disjoins primary claims from secondary claims. In addition, various sub-classes of primary claim appear to be implicitly disjoined. Hohfeldian claim-rights are one such sub-class. But, for example, Wallace also proposes that some [primary] claims are the correlatives of *imperfect* directed duties (pp. 202–205). These are presumably a distinct sub-class.

As I have already signaled, there are some bits of morality he is wholly content to ignore. As a guide, he describes the focus of his concern as the morality of "right and wrong" (p. 3).

One effective way to finesse this boundary question is to fix on Wallace's assumption that all duties are directed duties. Moral duties fall squarely within the morality of right and wrong. So if there are some undirected duties, not only will these duties lacking a direction be cast out of his relational nexus, but their ejection will not be safe for him to ignore either. Yet even then, one last complication continues to muddy the waters.

The difference between directed and undirected duties can be re-described as the difference between duties that correlate with a claim and duties that do not. This means that, for a given putatively undirected duty, the "cost" of denying that it lacks a direction is either having to deny that it is a duty in the first place or having to affirm that it *correlates with a claim* after all. For those who defend a particular theory of what a claim or a claim-right is,[18] the second of these ways out will have some bite to it. However, for Wallace, who both deploys a sui generis notion of a "claim," in place of a claim-right, as the correlative of a directed duty, and treats that same notion as a theoretical primitive, the second way out comes more or less for free. As we have already noted, Wallace is fully prepared to affirm that there are claims correlative to imperfect duties, which can be understood as a bid to integrate duties of virtue within his relational nexus.

Now, for myself, I am more than happy to uphold the commonsense counterexamples—your favorite virtues, say, along with duties to supply your favorite public goods.[19] Anyone who shares the commonsense view is welcome to join me in this, which further restricts the scope of the relational nexus. That is all to the good. However, for my second exhibit, I would prefer to minimize the potential for fruitless wrangling. Rather than focus on some particular ordinary duty, then, I shall take up the case of the all-things-considered ought, which I abbreviate as "ATC ought." The ATC ought expresses the final verdict of morality on how a given agent is required to act in the situation at hand.[20]

To begin with, I shall argue that the ATC ought has no direction.[21] It is not owed to anyone. In this respect, it resembles an undirected duty, even if one thinks there are no undirected duties. Equivalently, no one has any claim that correlates with the ATC ought. Of course, the content of someone's claim may *coincide with* what the agent ought ATC to do. Recall my opening promise to Mia, my neighbor, that I would paint her fence next week. On the basis of my promise, Mia has a claim against me that I paint her fence next week. If my correlative duty to Mia is not defeated by competing moral considerations, it will also turn out that I ought ATC to paint her fence next week. However, the fact that Mia's claim and what I ought ATC to do coincide in their content does not mean that her claim *correlates with* what I ought ATC to do. The correlative of Mia's claim is my (undefeated) directed duty to paint her fence and not "what I ought ATC to do"—even though, here, painting her fence *is* what I ought ATC to do. What an agent ought ATC to do is not, as such, something that figures inside the content of anyone's claim against them.

To bring out the pivotal difference between correlation and coincidence, consider what follows if Mia releases me from my promise to paint her fence (perhaps she subsequently caught a glimpse of my last paint job). It follows immediately that my undefeated directed duty to paint the fence is

[18]Like Wallace, Matthew Kramer (1998) insists that all duties are directed duties. Unlike Wallace, Kramer accepts that directed duties are analyzable, and he holds a version of the ordinary interest theory of rights. To defend his position, Kramer is prepared to ascribe a correlative claim-right in the face of any and every putatively undirected duty (1998, pp. 24–25).

[19]In the latter case, for those willing to analyze the correlation between claim-rights and directed duties, that is, to analyze (HCA), we can even explain why these duties are undirected. For discussion, see Sreenivasan (2017).

[20]Throughout I refer to the ATC *moral* ought. Nothing in the text implies that moral requirements are "overriding." Nor do I take any position on the question of how the moral ought is related to any generic practical ought.

[21]In his contribution to this symposium, Stephen Darwall argues for the same conclusion. As will be evident on inspection, however, Darwall and I argue for this shared conclusion on wholly different grounds (albeit entirely compatible grounds). See Darwall (2026).

dissolved. But even if it is no longer the case that I ought ATC to paint her fence, this does not follow immediately from Mia's release. Morality's final verdict on my action in this situation also depends on further moral facts, such as whether other things are equal or it is otherwise morally irrelevant whether I paint the fence. For example, suppose that, in the interim, a proverbial evil demon or criminal mastermind threatens to wreak havoc if I do not paint Mia's fence. It will then continue to be the case that I ought ATC to paint it, notwithstanding her release. Thus, unlike my promissory duty to Mia, the ATC ought is not controlled by her power to release me from my promise. Unlike my promissory duty to Mia, the ATC ought does not correlate with her claim.

It remains to license the classification of the ATC ought as a "duty." Or, what comes to the same thing, to license the inference from "the ATC ought has no direction" to "some duties have no direction." In the first instance, this is primarily a matter of how fussy one is about terminology. If one is not at all fussy, the step can be taken immediately. Ultimately, however, the terminological question is neither very interesting nor very helpful to Wallace.

Wallace has a nice discussion of the ways in which having a "duty" to φ differs from (merely) having a "reason" to φ (pp. 126–130), eloquently summarizing the conventional philosophical wisdom about the significance of these differences. In brief, duties have a certain peremptory force that reasons lack, but which is characteristic of the distinctive quality exhibited by moral demands. Wallace groups "obligation" and "requirement" together with "duty," on the favored side of this dividing line. He suggests that they can all be regarded as forms of "constraint," though reasons are not constraints. One question, then, is whether the moral "ought" can also be regarded as a constraint in just this sense. (I do not see why not.)

However this matter is resolved, it cannot be denied that *the ATC moral* ought is a constraint in Wallace's sense. It has at least as much moral "force" as a directed duty.[22] Or perhaps better, since the feature at issue is not really a matter of quantity, the ATC moral ought clearly has *conclusive* force.[23] After all, it expresses the final verdict of morality in the case at hand. This places the ATC moral ought fully on a par with directed duties, as far as the dimension of force is concerned.

In the present context, the dispute about whether there are undirected duties is fundamentally a dispute about whether the two features joined in a directed duty—having a direction and being a constraint—can also come apart. On reflection, then, the fact that the ATC ought has no direction is already enough to demonstrate that the two features do come apart. Moreover, they come apart whether or not the ATC ought is classified as a "duty." That is all we need to show.

In fact, if we take a few steps back, it will become apparent that we did not even need to show that much. For the reason we fixed upon "duties" in the first place was that they were a reliable proxy for elements of morality that fell squarely within what Wallace calls the "morality of right and wrong." However, the ATC moral ought qualifies immediately for that title, without any need for a proxy. And yet, despite belonging squarely to the morality of right and wrong, the ATC ought is still excluded from the relational nexus, on account of its lacking any direction. I therefore offer the ATC ought as exhibit two in my catalog of where the scope of the relational nexus fails to reach Wallace's ambition.

3. Directed Duties are Not Absolute

My critique of Wallace's portrait of morality as a relational nexus runs along two axes. One axis concerns the scope of the relational nexus, while the other axis concerns two extraneous features that Wallace ascribes to directed duties. My argument against his first extraneous feature—directed duties as unanalyzable primitives—has already been given (at the end of Section 1). I shall now argue against his thesis that directed duties are absolute.

[22]The ATC ought obviously has more moral force than a pro tanto duty. As we review in Section 3, though, Wallace understands directed duties to be conclusive duties, rather than pro tanto duties. While I reject this understanding, my formulation aims to be ecumenical.

[23]It is therefore important not to confuse the ATC moral ought with the "balance of reasons."

To set the stage, recall that an absolute duty to φ is a duty to φ that *entails* that its bearer ought ATC to φ. Wallace sometimes expresses the same concept by saying that directed duties are "conclusive." In either formulation, it is important to be clear that particular directed duties cannot be absolute or conclusive on their own, as distinct from other directed duties. Either all directed duties are absolute or none is. (Likewise for their being conclusive.) Thus, Wallace's position is that all directed duties are absolute, whereas mine is that no directed duties are absolute.[24]

If directed duties are absolute, they will function a little differently from what may seem intuitive. In particular, since they cannot come apart from the ATC ought, the final verdict of morality on the situation at hand always has to be settled *before* any directed duty can hold there. Otherwise, the entailment from a given directed duty would risk being false. For example, on Wallace's view, promisors do not immediately have a directed duty to perform on their promise. Their coming to have a directed duty to perform depends on whether, in the event, they ought ATC to perform. If it turns out that, ATC, a promisor ought *not* to perform on their promise, it is not that their promissory duty is defeated, but rather that they never had any such duty in the first place, notwithstanding the validity of their promise.

One way to describe how directed duties function on this view, then, is that they are *outputs from* moral deliberation or evaluation, rather than inputs to it. Claims function similarly (pp. 170–171). Whether this description represents an objection to the view, or simply something to get used to, is one of the bones of contention in a long-running debate (e.g., between Thomson, 1990 and Wallace, 2026). As I suggested earlier, I shall not revisit the previous debate. Instead, I shall argue that Wallace's position should be rejected on entirely different grounds. However, to avoid prejudicing the terrain against him, I shall make every effort to accommodate the terms of my discussion to the logic of his position. Since this may give some of the discussion a counter-intuitive flavor, I wanted to explain clearly up front whence it comes and why.

Considerable leverage can be applied against the thesis that all directed duties are absolute by taking due account of how it intersects the debate between consequentialism and deontology. Shelly Kagan (1998) usefully characterizes the difference between two species of deontology in terms of a distinction between two kinds of "deontological constraint." As Kagan defines it, a *deontological constraint* against φ-ing entails that φ-ing is ATC impermissible, even when φ-ing is necessary to produce the most overall good in the situation at hand (1998, p. 72). For example, a deontological constraint against killing entails that killing one to save five is ATC impermissible, even when killing the one is the *only* way to save five other lives (and saving the five represents the most overall good that can be produced under the circumstances). A salient, and intentional, feature of Kagan's definition is that the existence of deontological constraints falsifies consequentialism, since under consequentialism it is never impermissible to produce the most overall good.

Deontological constraints can be either "moderate" or "absolute." They are *absolute* when there is no restriction on the magnitude of overall good that is irrelevant to the entailed impermissibility of φ-ing. Any and all magnitudes of overall good are irrelevant (even when φ-ing is necessary to produce that magnitude of good). For example, an absolute deontological constraint against killing entails that killing one person is ATC impermissible, even when that is necessary to save, say, a billion others. Since the generic definition of a deontological constraint given earlier includes no such restrictions, what it describes, strictly speaking, is an absolute constraint.

By contrast, *moderate* deontological constraints incorporate what Kagan calls a "threshold," which restricts the magnitude of overall good that is irrelevant to the entailed impermissibility of φ-ing (1998, p. 79). Below the specified threshold, moderate constraints operate similarly to

[24]If no directed duties are absolute, it can still easily turn out, for a particular directed duty to φ, that its bearer ought ATC to φ. In other words, a given directed duty can easily wind up coinciding with (or being confirmed by) the ATC ought. Among other things, this will happen whenever an agent has a directed duty and other things are equal.

absolute ones.[25] However, when the magnitude of overall good that can be produced by φ-ing *exceeds* the specified threshold, and φ-ing is necessary to produce it, moderate deontological constraints become consistent with the ATC permissibility of φ-ing. Under those conditions, the distinctive entailment of a deontological constraint no longer holds, even though the deontological constraint itself remains in place. Finally, with moderate constraints, the specified threshold can vary from constraint to constraint. For example, the moderate deontological constraint against killing can have a higher threshold than the moderate deontological constraint against assault has.

Against this background, "absolutist" deontology can be distinguished as the species of deontology holding that some deontological constraints exist and all of them are absolute constraints. "Moderate" deontology, in turn, holds that some deontological constraints exist and some of them are moderate constraints.[26] As this typology reminds us, then, the debate between consequentialism and deontology is actually at least a three way debate.[27] Absolutist and moderate deontologists disagree with each other about whether any deontological constraints have thresholds, while agreeing that some deontological constraints exist (and hence, that consequentialism is false). Moderate deontologists can accept that there are some absolute constraints, but they need not affirm it.

To establish the fulcrum that will enable leverage to be applied against Wallace's position, observe that if all directed duties are absolute, my having a directed duty not to assault Meredith (say) is *equivalent* to the existence of an absolute deontological constraint against assault, at least so far as Meredith is concerned.[28] For then my directed duty entails that I ought ATC not to assault Meredith, that is, that it is ATC impermissible for me to assault her. Moreover, and crucially, this entailment is *not conditioned* on the magnitude of overall good I can produce "by" assaulting Meredith (i.e., can produce if and only if I assault her). It is the absence of any such conditioning that makes the entailment from my directed duty not to assault Meredith equivalent to the entailment from an absolute deontological constraint against assault.

Now, to apply the leverage. Let us assume, for the sake of a reductio, that moderate deontology captures the truth about morality. For concreteness, suppose that any deontological constraint against *assault* incorporates a suitable threshold. Unfortunately for Wallace, it follows that Meredith has no claim against me that I not assault her. The problem is not that Meredith has no such claim against me initially. Recall that, on Wallace's view, no one enters any situation with any claims against others. Claims are outputs, he believes. Rather, the absurdity is that Meredith can never end up with a claim against me that I not assault her either. For if she had such a claim against me, it would entail that I had a directed duty not to assault her. As we have seen, that duty would be equivalent to an absolute deontological constraint against assault (cf. note 28). Yet, ex hypothesi, there is no absolute deontological constraint against assault. Hence, I can never have a directed duty not to assault Meredith, and she can never hold the correlative claim against me.

---

[25]Below its threshold, a moderate constraint against φ-ing also entails that φ-ing is ATC impermissible, even when φ-ing is necessary to produce the most overall good. Hence, the existence of moderate deontological constraints falsifies consequentialism, too. However, there is also a vital dissimilarity, insofar as a moderate constraint's entailment is *conditioned* on the magnitude of overall good in play being below the threshold.

[26]Wallace (2019) discusses moderate deontology and thresholds in his Chapter 6. I engage what he says there in Section 4 below.

[27]For simplicity, I ignore the possibility of divisions within consequentialism.

[28]For yet more simplicity, I shall abstract from the fact that deontological constraints have universal application, whereas particular directed duties do not. (That is, an absolute deontological constraint against assault entails not only that it is ATC impermissible for me to assault Meredith, but also that it is ATC impermissible for me to assault, for example, Mabel or Michael. By contrast, the duty I owe Meredith not to assault her does not entail that I have a duty not to assault Michael.) In some contexts, this is an important difference, but not the present one. However, I will not always repeat the qualification in the text that makes this abstraction explicit.

It may be helpful to bring out the fundamental inconsistency between deontological constraints with thresholds and Wallace's conception of directed duties in a slightly different way, one that avoids any use of the labels "absolute" or "moderate." If some domain of morality—assaults, say—is governed by a deontological constraint that incorporates a threshold, then nothing can license any conclusion about the ATC permissibility of a given assault *without* relying on some premise locating the magnitude of good that can be produced by that assault in relation to the threshold governing assaults (cf. note 25).[29] As Wallace conceives directed duties, however, and therefore also their correlative claims, my directed duty not to assault Meredith immediately entails that it is ATC impermissible for me to assault her—that is, it entails this without any reference to the magnitude of good I can produce by assaulting her. But if nothing can have feature F, and Wallace-approved claims against assault have feature F by definition, then Wallace-approved claims against assault do not exist. It follows that Meredith has no such claim against me.

Of course, claims not to be assaulted are merely an illustration. What actually follows from Wallace's position is much worse: No one can have any claim against anyone in *any* domain of morality governed by a moderate deontological constraint. But claims that others not φ, for various everyday values of "φ," should not evanesce just because there are some scenarios in which it would nevertheless be ATC permissible for someone to φ. That is absurd. However, it is a consequence of the thesis that all directed duties are absolute (under moderate deontology). We should therefore reject that thesis.

But why accept the premise of my argument? Might moderate deontology not be false? Indeed, it might be. To begin with, however, I should clarify that a much weaker premise would serve me just as well. Namely, either moderate deontology is true or consequentialism is true.[30] I focused on the first disjunct purely for rhetorical purposes, since moderate deontologists are much more likely to be sympathetic to Wallace's portrait of morality than consequentialists are.

For convenience, we can divide the dialectical universe into two camps, committed partisans of one of the three positions in the debate and uncommitted external observers. Evidently, my original premise is music to the ears of committed moderate deontologists. With committed absolutist deontologists, I accept that my argument exerts no leverage. In effect, their position already commits them to Wallace's thesis anyhow. That leaves the external observers (cf. note 30). They face a choice among committing themselves to absolutist deontology (the least plausible position in the entire debate), rejecting the thesis that all directed duties are absolute, or putting all manner of everyday claims against others at grave risk of evanescence. In the absence of a convincing argument for absolutist deontology, which Wallace does not attempt, this is a completely straightforward choice. Apart from committed absolutist deontologists, everyone has good reason to reject the thesis that all directed duties are absolute.

Thus far in this section, I have been completing the second axis of my critique of Wallace. Yet there is also a further fruitful point of interaction between the two axes to register. On the one hand, as we have seen, the mechanism Wallace employs to expand the scope of his relational nexus is assimilation to the case of promising. On the other hand, as we have just been discussing, he affirms that all directed duties are absolute. This means that, on his view, promissory duties, the target of any effort at assimilation, are absolute duties. However, and finally, moderate deontological constraints are inconsistent with the thesis that all directed duties are absolute. So we have just concluded. It follows that they cannot be assimilated to promissory duties *as Wallace conceives them*, and are consequently excluded from his relational nexus. We can therefore add moderate deontological constraints, all of them, to our running catalog of where Wallace's nexus falls short of his ambition. They are my third exhibit and our final nettle.

[29]For these purposes, it does not matter whether the assault in question is ATC permissible or ATC impermissible.

[30]An alternative argument can be constructed in parallel to mine, beginning with the premise that consequentialism is true and ending in results equally, if not more, uncongenial to Wallace.

4. Objections and Replies

Let me close the article by entertaining two objections to what I have said. The first objection picks up from where we just left off. It complains that I have been somewhat unfair to Wallace. In my view, the thesis that all directed duties are absolute is both false and extraneous to the relational character of promissory duties. But I still let the fact that moderate deontological constraints are inconsistent with this extraneous feature suffice to expel them from Wallace's relational nexus. What if Wallace were to abandon this thesis instead? Would that not enable him to reclaim moderate deontological constraints for his nexus, without affecting its relational character in the least? Should that option not be considered?

I accept part of what this objector believes. Wallace can, of course, abandon the false thesis, and he certainly should. The objector is also correct that this move would not alter the relational character of Wallace's nexus. In that sense, it represents a pure improvement in his position. What is more, abandoning the thesis that all directed duties are absolute is plainly a *necessary* condition on reclaiming moderate deontological constraints for the relational nexus. However, it is not a sufficient condition. It would not itself enable them to be reclaimed. Let me explain.

The main point is that moderate deontological constraints do not *have to be* understood or structured as directed duties. Indeed, the most natural way to understand them is in terms of rules, that is, as simple prohibitions. Do not kill; do not assault; and so on. A rule against φ-ing must have two features in order to play the role of a moderate constraint against φ-ing, both of which will be familiar from our previous discussion. First, the rule must sometimes make it ATC impermissible to φ, even when φ-ing is necessary to produce the most overall good in the situation at hand.[31] Second, the rule must also have a "threshold," something that limits the first feature (or turns it off), so that at other times φ-ing is ATC permissible.[32] The higher the threshold, the more often or widely will it be ATC impermissible to φ.

To vindicate the existence of a moderate deontological constraint, one has to establish that "something" has both of these features (or, at least, that both features are really at work), for some value of "φ." Whether any such argument has ever been produced is one of the central questions driving the debate between consequentialism and deontology. However, *if* an argument of this kind succeeds, its conclusion can be reported without loss in terms of a rule's having the two features in question—or, perhaps better, in terms of a rule's having the two corresponding effects on the ATC ought. And yet rules themselves have no relational character at all. They are impersonal. Nor, as we have seen, does the ATC ought. The vindication of moderate deontology does not, therefore, require anything with a relational character to be invoked. So far as moderate deontological constraints are concerned, a relational character is actually the extraneous feature.

None of this categorically closes the door on making the case that, even though a relational character is entirely unnecessary for their operation, moderate deontological constraints are nevertheless best understood, all in all, in terms that centrally incorporate directed duties. It is hard to see how this could succeed, but it is not impossible. All the same, until a convincing argument to this effect is in hand, the simpler interpretation is clearly better. On the simpler interpretation, moderate deontological constraints have no relational character. Hence, if they exist, they do not belong in any relational nexus.

The second objection takes issue with my conclusion that the ATC ought has no direction. In summarizing the dialectical context of arguments over the existence of undirected duties, in Section 2, I described two paths for rejecting would-be counterexamples to the thesis that all duties

[31]Exactly how a rule brings this about is not something we need to decide. One way to have this effect would be to assign sufficient negative moral weight to φ-ing, but it is not the only way.

[32]In the absence of this second feature, a rule with (only) the first feature plays the role of an absolute deontological constraint.

are directed duties: either denying that the example on offer is a duty or affirming that it has a correlative claim. With respect to the ATC ought, the second objector takes the first path, but goes one better, denying that the ATC ought even exists.[33]

On the objector's picture, the ATC ought is really only the shadow, as it were, of the last claim standing in the situation at hand. Or, more precisely, since claims are outputs in Wallace's view, it is the shadow of the only claim to stand there.[34] When I have made conflicting promises, for example, at least two potential claims against me will be in prospect. Moral reasoning will be required to work out which of my promisees actually holds a claim to my performance. Once that is settled, morality requires me to perform the promise made to the promisee holding the claim. That is, it requires me to act on my duty *to that* promisee. If we like, we can add that this directed duty entails that I ought ATC to perform that promise. That is not incorrect. But it does not add anything of substance to the analysis either. In this context, the ATC ought is not materially distinct from my directed duty. It only serves to put a certain formal gloss on my duty; and likewise with other directed duties in other contexts. Hence, even if the ATC ought is "excluded" from the relational nexus, nothing of moral substance is thereby missing from Wallace's portrait. For present purposes, that makes it a rather poor exhibit.

Roughly speaking, this objection has two halves. In its first half, the objector contends that the content of the ATC ought always coincides with someone's claim. In its second half, he or she then infers that the ATC ought is a purely formal epiphenomenon, and so can be dismissed, since in any given case it is not materially distinct from some directed duty. I am not sure we should accept the objection's distinctive inference. But it does not really matter. For the inference presupposes the set up proposed in the first half of the objection and we can reject that setup.

Extensionally, the first half of the objection is almost equivalent to Wallace's thesis that all duties are directed duties. It is precisely equivalent to the conjunction of that thesis and the proposition that only duties can determine the ATC ought. Recall that, in our original discussion of Wallace's thesis, I mentioned several commonsense counterexamples, such as duties to supply various public goods. It bears repeating that accepting any such counterexample suffices to reject his thesis, and hence this second objection. I did not insist on common sense then, in order to avoid fruitless wrangling. Throughout the paper, I have also been trying to avoid the gray zone in between the outer edges of Wallace's ambitions for the relational nexus and the border of morality proper. However, the time has now come when we shall have to discard these strictures.

In the last chapter of his book, Wallace considers a sequence of cases that can, even he admits, seem "especially resistant to treatment with a relational framework" (p. 191). For the most part, Wallace holds his ground down this line, arguing that the challenging cases can actually be accommodated, although he concedes they cannot all be accommodated. To show where the objection goes wrong, let me take up two of the cases he discusses, one more briefly and brusquely than the other.

I shall begin with duties to future generations. Consider, for example, a duty to reduce current consumption (perhaps, drastically), so as to mitigate the effects of current consumption on continuing environmental degradation and climate change. Is this a directed duty? Is it a duty we owe *to* individual members of future generations? Wallace wrestles with a familiar impediment to an affirmative answer here, arising from the non-identity problem (pp. 210–215). Ultimately, he contends that recourse to the device of "non-comparative interests" enables claims to be ascribed to future individuals—claims that correlate with our duty to reduce current consumption—despite

[33]I am grateful to Sergio Tenenbaum for raising this objection. Cf. Wallace (2019, p. 55).

[34]For simplicity, I begin by ignoring cases that might best be described as resulting in an irresolvable moral dilemma. Dilemmas will make an appearance later.

the fact that different future individuals will exist in the scenario where we comply with our duty, as compared to the scenario in which we do not.

Now the non-identity problem is notoriously subtle and vexed; and I shall not attempt to evaluate Wallace's effort to extricate himself from it. But considering how absorbed he is by the non-identity problem, it is somewhat surprising that Wallace pays no attention at all to a much simpler problem, which might be called the non-existence problem (de George, 1981). Future people do not exist. They will exist in the future (if the future comes). But they do not exist at the same time that we exist. If existence is the minimum prerequisite for claim-holding, as seems entirely reasonable, then future people cannot hold any claims against us. On this point, it does not matter how many different sub-classes of claim one is prepared to distinguish. Their minimum prerequisite fails all the same. Hence our duty to reduce current consumption, which I am perfectly happy to accept, cannot be a duty we owe to individual members of future generations. That makes it an excellent candidate for being an undirected duty.

The last kind of case in Wallace's sequence is one where the philosophical question is whether, to use the classic phrase, the numbers count (pp. 221–233). More specifically, in these cases, one side of the agent's choice will inflict a small loss (or a small risk of a large loss) on many people, while the flip side of the choice will inflict a large loss—most simply, (certain) death—on one person.[35] In this context, the challenge for Wallace is whether and how to accept that there is some number of people, given which the aggregate loss on the many people side of the choice is sufficient to *counteract the claim* held by the singleton not to suffer the large loss. As he acknowledges, many philosophers have the strong intuition that there is some such number. Eventually, Wallace does concede that "considerations of aggregate well-being" have enough "normative significance" for "practical thought" to enable them to play the role represented by that number (pp. 227–230).

Even then, however, Wallace enters two caveats, which are meant to insulate the relational nexus from his concession. According to the first caveat, what follows from the relevant claim's being counteracted is *not* that it is ATC permissible to inflict the large loss on the singleton. Rather, what follows is that the agent in this situation faces a "practical dilemma" (p. 231). That is, the agent is subject to a conflict between two independent practical requirements, neither of which defeats or is defeated by the other. According to the second caveat, the requirement to inflict the large loss on the singleton, which is generated by considerations of aggregate well-being, is *not a moral* requirement, but instead an "extra-moral" requirement. Thus, the practical dilemma takes the form of "a conflict between moral and extra-moral requirements" (232).

To grasp the significance of these caveats, it helps to contrast the form of pluralism Wallace is prepared to concede with another form, which he is straining to avoid. While Wallace associates it with one of his foils (Frick, 2015), we can call it commonsense pluralism. On commonsense pluralism, the practical requirement generated by considerations of aggregate well-being is also a moral requirement. Hence, the conflict in cases where the numbers count is a conflict entirely *within* morality.[36] Furthermore, the conflict need not result in a dilemma. At least when the loss to aggregate well-being is big enough, commonsense pluralism can judge that, ATC, the agent is morally required to avoid it (and so, required to inflict the large loss on the singleton).

As should be manifest, commonsense pluralism contradicts the setup in the first half of the objection we are examining. On this view, it is simply false that the content of the ATC moral ought always coincides with someone's claim. After all, in some cases where the numbers count, no claims are held against the agent at all and the ATC moral ought still requires the agent to inflict the large loss on the singleton. Its content in these cases is therefore materially distinct from that of any directed duty.

[35]Wallace conducts this discussion with reference to two examples, one that involves risk (Frick 2015, pp. 212–219) and another that does not (Parfit 2011, p. 197).

[36]Indeed, in many respects, it resembles a conflict of promises, as described by Wallace and the second objector.

We can confirm this description by inspecting both sides of the stylized choice facing the agent in these cases. On the side of the one, the singleton has no claim, but that need not be because their claim has been defeated. The situation can also be understood, following Wallace, as one where the singleton's claim has ultimately been prevented from arising—just like the claim of the losing promisee in a conflict of promises, except here the prevention is effected by having aggregated considerations of well-being. On the side of the many, there need not be any claims either. Even when aggregate well-being generates a moral requirement on the agent, the increments of well-being that get aggregated need not be increments to which any of the many has a claim on their own; and aggregation across individuals never brings an individual claim into existence (out of thin air, as it were).[37]

Accordingly, the question we must confront is whether the caveats to the pluralism Wallace concedes do anything to stave off the clear implication that emerges under commonsense pluralism: namely, in some cases where the numbers count, the ATC moral ought is materially distinct from any directed duty, and so cannot be dismissed as a purely formal epiphenomenon. As I shall now explain, they do not. I conclude that we can reject the second objection.

As Wallace himself recognizes (p. 190), the primary disagreement between the two forms of pluralism here concerns the correct answer to a first-order moral question. We can frame it as the question whether—in cases where the numbers count, and the numbers are there—it is ATC morally permissible to inflict the large loss on the singleton.[38] In principle, however, any of the possible answers to this first-order question might be correct. But this means that, in its nature, the concept of the ATC moral ought has to be able to accommodate the full range of possible answers, including commonsense pluralism's. Since this accommodation could not be made by a purely formal epiphenomenon, the ATC moral ought cannot be purely formal.

To put the same point somewhat differently, if the ATC moral ought necessarily coincided with some directed duty, then commonsense pluralism's answer would be conceptually incoherent. But that is much too strong. We could turn this into a reductio of the second objection if we had to. Except there is no need, since Wallace himself explicitly accepts that commonsense pluralism's answer is possible (p. 230). He thinks its answer is incorrect, not incoherent. Hence the ATC moral ought does not necessarily coincide with some directed duty, and so the second objection fails.[39]

The objections we have examined aim to remove one exhibit each from our catalog of limitations on Wallace's ambition to expand his relational nexus outward from promising over the rest of morality. Since we have rejected both objections, moderate deontological constraints and the ATC moral ought remain secure in their status as exhibits in this catalog.[40] Neither belongs in the relational nexus.

---

[37]Notice that by Wallace's own admission, then, a practical requirement (i.e., a duty) need not, as a conceptual matter, correlate with a claim. In other words, conceptually, a generic duty need not be a directed duty. Wallace still refuses these generic undirected duties the title of "moral duty," as we have seen. But the conceptual point about duties per se stands independently of their labeling.

[38]In a dilemma, none of the options is permissible. So on Wallace's view, it is never morally permissible to inflict the large loss on the singleton (at least, not in these cases).

[39]In discussing the second objection, I have narrowly tailored my interpretation of it to the requirement of serving Wallace's purposes. That is why I took the proposition that the content of the ATC ought always coincides with someone's claim as my point of departure. However, some advocates of a purely formal analysis of the ATC ought will have no interest in Wallace's project. These advocates are by no means committed to my point of departure. They can fully accept the existence of undirected duties or, indeed, aggregate well-being as a free-standing and non-relational moral consideration, and nevertheless insist that the ATC ought is still to be analyzed purely formally. We have not explored the grounds on which this somewhat different position might be defended, since it is plainly anathema to Wallace. At the same time, nothing in my rebuttal of the second objection excludes such an anti-Wallace formalism either. For present purposes, its merits can be left open.

[40]Anti-Wallace formalists from note 39 can replace the ATC moral ought with (say) aggregate well-being as a non-relational moral good and still keep hold of three significant elements of morality that resist assimilation to the relational nexus.

5. Conclusion

Wallace portrays morality as a nexus of relational features, taking inspiration from the central case of promising. Notably, a promisor has a directed duty to the promisee, who in turn holds a correlative right against the promisor. Wallace expands Hohfeld's classical conception of the correlativity present in promising, from (HCA) to (W), and also proliferates the number of "claims" he is prepared to recognize. Wielding this apparatus, he attempts to assimilate great swathes of morality to the promissory paradigm.

Against Wallace's project, I objected that significant elements of morality are inevitably excluded from his relational nexus. They are either overly different from the case of promising or they lack a relational character altogether. My three exhibits on this front were relationship-based wrongings, the ATC ought, and moderate deontological constraints. By sharply curtailing the grandeur of his ambition for the relational interpretation of morality, Wallace can do better justice to the rich insight of his original vision.[41]

Gopal Sreenivasan is the Lester Crown University Distinguished Professor of Ethics at Duke University. He is the author of *Emotion and virtue* (Princeton, 2020) and the co-editor, with Margaret Gilbert and Jeffrey Helmreich, of the *Palgrave Handbook on the Philosophy of Rights* (Palgrave Macmillan 2026).

References

Cornell, N. (2015). Wrongs, rights, and third parties. *Philosophy & Public Affairs*, 43(2), 109–143. https://doi.org/10.1111/papa.12054

Cornell, N. (2025). *Wrongs and rights come apart*. Harvard University Press. https://doi.org/10.2307/jj.21275861

Darwall, S. (2026). Why obligations can't be relational all the way down. *Canadian Journal of Philosophy*, 1–11. https://doi.org/10.1017/can.2026.10099

de George, R. (1981). The environment, rights, and future generations. In E. Partridge (Ed.), *Responsibilities to future generations: environmental ethics* (pp. 157–166). Prometheus Books.

Frick, J. (2015). Contractualism and social risk. *Philosophy & Public Affairs*, 43(3), 175–223. https://doi.org/10.1111/papa.12058

Hohfeld, W. (1919). *Fundamental legal conceptions, as applied in judicial reasoning*, ed. W. Cook. Yale University Press.

Kagan, S. (1998). *Normative ethics*. Westview Press.

Kramer, M. (1998). Rights without trimmings. In M. Kramer, N. Simmonds, & H. Steiner (Eds.), *A debate over rights* (pp. 7–111). Clarendon Press.

Martin, A. (2021). Personal bonds: Directed obligations without rights. *Philosophy and Phenomenological Research*, 102(1), 65–86. https://doi.org/10.1111/phpr.12620

Parfit, D. (2011). *On what matters* (Vol. 2). Oxford University Press. https://doi.org/10.1093/acprof:osobl/9780199572809.001.0001

Sreenivasan, G. (2005). A hybrid theory of claim-rights. *Oxford Journal of Legal Studies*, 25(2), 257–274. https://doi.org/10.1093/ojls/gqi013

Sreenivasan, G. (2017). Public goods, individual rights, and third party benefits. In M. McBride (Ed.), *New Essays on the Nature of Rights* (pp. 127–148). Hart Publishing.

Sreenivasan, G. (2024). Rights against the world. *Analysis*, 84(2), 311–319. https://doi.org/10.1093/analys/anad067

Thomson, J. (1990). *The realm of rights*. Harvard University Press.

Wallace, J. (2019). *The moral nexus*. Princeton University Press.

Wallace, J. (2026). Relational obligations and the status of moral claim rights. In M. Gilbert, J. Helmreich, & G. Sreenivasan (Eds.), *Palgrave Handbook on the Philosophy of Rights*. Palgrave MacMillan.

[41]For helpful comments on a previous version of this paper, I am very grateful to Rowan Cruft, Tom Dougherty, Sergio Tenenbaum, and Jay Wallace. Many thanks to Rahul Kumar and Jonas Vandieken for inviting me to participate in their workshop on relational morality, held at the Centre for Ethics at the University of Toronto; and further thanks to them and the other participants for helpful discussion at the workshop.

Cite this article: Sreenivasan, G. 2026. Nettles in the Nexus: On Wallace on Relational Morality. *Canadian Journal of Philosophy* 56: e24, 1–17, doi:10.1017/can.2026.10073