

# Two concepts of directed obligation

Brendan de Kenessey 

University of Toronto, Toronto, Canada

## Correspondence

Brendan de Kenessey, University of  
Toronto, Toronto, Canada.

Email: [brendan.dekenessey@utoronto.ca](mailto:brendan.dekenessey@utoronto.ca)

## Funding information

Connaught Fund, Grant/Award Number:  
Connaught New Researcher Award

## Abstract

This paper argues that there are two importantly distinct normative relations that can be referred to using phrases like ‘X is obligated to Y,’ ‘Y has a right against X,’ or ‘X wronged Y.’ When we say that I am obligated to you not to read your diary, one thing we might mean is that I am subject to a deontological constraint against reading your diary that gives me a non-instrumental, agent-relative reason not to do so, and which you are typically in a unique position to waive with consent. I call this first relation *the constraint relation*. A second thing we might mean is that you are in a position to fittingly hold me personally accountable for reading your diary by demanding that I not read your diary, resenting me if I do so without excuse, and deciding whether to forgive me for this afterwards. I call this second relation *the accountability relation*. Though these two kinds of directed obligation often coincide, I argue that they are extensionally dissociable and play different normative roles. We cannot provide an adequate theory of ‘obligation to’ until we recognize that this phrase denotes not one relation, but two.

## KEYWORDS

accountability, blame, constraints, directed duties, directed obligation, inalienable rights, rights, specificationism

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2024 The Author(s). *Philosophy and Phenomenological Research* published by Wiley Periodicals LLC on behalf of Philosophy and Phenomenological Research LLC.

Many of our obligations are owed to particular people. I am not just obligated to call my mother on her birthday; I am obligated *to my mother* to do so. If I stomp unbidden on your sand castle in a fit of destructive glee, then I have not merely done wrong: I have wronged *you*.

There has been a recent surge of interest in this phenomenon, which I will call *directed obligation*.<sup>1</sup> Yet this exciting literature has mostly failed to recognize that there are in fact two distinct phenomena picked out by our everyday talk of obligation to others. Philosophers have so far treated directed obligation as a single thing, while in fact there are two distinct relations both worthy of the name. This paper's aim is to pull apart these two concepts of directed obligation.<sup>2</sup>

Briefly, the distinction is this. One relation, which I call *the constraint relation*, is the relation I stand in to you when you have certain paradigmatic rights against me, such as promissory rights, bodily rights, property rights, and privacy rights. When you have a right of this kind that I  $\phi$ , I am subject to a deontological constraint against not- $\phi$ ing that gives me a non-instrumental, agent-relative reason to  $\phi$ . I will call the directed obligations implied by this relation *constraint obligations*, and the rights that are the flip side of these obligations *constraint rights*. The second relation, which I call *the accountability relation*, is the relation I stand in to you when you have warrant to hold me personally accountable for whether I  $\phi$  by demanding that I  $\phi$ , resenting me if I fail to  $\phi$  without excuse, and deciding whether to forgive me *for  $\phi$ ing*. I will call the directed obligations implied by this relation *accountability obligations*, and the rights that are the flip side of these obligations *accountability rights*.

These relations often go together: when I have a constraint obligation to you, that usually gives you the warrant to demand that I respect this obligation and resent me if I don't, thus implying an accountability obligation. But I will argue that these two relations are extensionally dissociable and play different roles in our moral thought.

The most important difference between the relations lies in their normative roles. As I will put it, the constraint relation is a normative *input*, while the accountability relation is a normative *output*. The fact that you have a constraint right that I  $\phi$  serves as a contributory input to the question of what I ought to do: your constraint right adds a normative reason in favor of my  $\phi$ ing to the overall balance of my reasons. The reason this constraint obligation provides, while strong, could be outweighed by other, stronger reasons for me not to  $\phi$ . In contrast, your demand that I  $\phi$  or resentment at my failure to  $\phi$  cannot be fitting unless I ought to  $\phi$  all things considered. Thus the fact that it is fitting for you to demand that I  $\phi$  or resent me for not- $\phi$ ing does not contribute to the balance of my reasons, but depends upon it. Whether I have an accountability obligation to you is determined by the output of my reasons as a whole.

This difference in role explains the concepts' different extensions. I can be under a constraint obligation to you without being under a corresponding accountability obligation because constraint obligations can be *overridden*: I might have decisive reason, all things considered, to

<sup>1</sup> Examples of this recent literature include Sreenivasan (2010), Hedahl (2013), Cornell (2015), Ripstein (2016), Gilbert (2018), Wallace (2019), and Martin (2021). Some philosophers prefer the term 'bipolar obligation': e.g., Darwall (2013) and Vandieken (2019).

<sup>2</sup> This phrase, and the title of this paper, is meant to echo Rawls's "Two Concepts of Rules" (Rawls, 1955). However, I will mostly speak of two *relations* rather than two 'concepts', as my claim isn't about human psychology, but instead about normative reality.

infringe your constraint right.<sup>3</sup> Conversely, I can be under an accountability obligation to you to  $\phi$  without being under a corresponding constraint obligation because non-deontological normative reasons can make personal accountability fitting. Specifically, I will argue that considerations of *welfare* can ground accountability obligations without the help of any constraint rights. In short, the constraint input can fail to generate an accountability output, and the accountability output can be produced by non-constraint inputs.

Some philosophers have already noticed this distinction, and there are distinct literatures that focus on each of our two relations.<sup>4</sup> But these two kinds of directed obligation are not distinguished clearly, explicitly, or frequently enough. If we don't pay more attention to this distinction, we are in danger of blending the two relations together. Firmly distinguishing between the accountability and constraint relations will allow us to develop better theories of both.

I am convinced that all of the ordinary language phrases theorists use to refer to directed obligation – 'obligation to,' 'duty to,' 'owed to,' 'wronging,' 'having a right,' 'having a claim,' 'being entitled to' – are ambiguous between the constraint and accountability senses. We can, and do, use any of these terms to refer to either normative relation. Thus we are in a 'jade'-like situation: we have one set of ordinary language terms for two importantly different phenomena.<sup>5</sup> The best way to avoid misunderstanding is to introduce technical terms that are stipulated to pick out only one concept. This is why I use the terms 'constraint relation' and 'accountability relation'.<sup>6</sup> I will also avail myself of the other phrases commonly used to denote directed obligation, appending the label 'constraint' or 'accountability' to mark which relation I mean. So I might say "X wrongs Y in the accountability sense" or "Y has a constraint right against X."

The rest of the paper proceeds as follows. §1 delineates the characteristic features of the constraint and accountability relations in more detail. §2 argues that there can be constraint obligations without corresponding accountability obligations. Along the way, I argue that the debate between 'specificationists' and 'generalists' can be resolved by appeal to this paper's distinction: specificationism is true for accountability obligations, generalism for constraint obligations. §3 argues that there can be accountability obligations without corresponding constraint obligations. §4 illustrates the usefulness of the paper's distinction by showing how it clarifies the question of whether there are any rights that cannot be waived by consent.

<sup>3</sup> Here I am following Judith Jarvis Thomson's terminology, using 'infringement' to refer to permissibly acting against another's rights, and 'violation' to refer to wrongfully acting against another's rights (Thomson, 1980: 122).

<sup>4</sup> Philosophers who have noticed this paper's distinction, or something close to it, include Cornell (2015), Ripstein (2016: 14–18), Wallace (2019: 2), Martin (2021), and Attie-Picker (2023). For examples of the literature on the constraint relation, see Hohfeld (1913), Hart (1955), MacCormick (1977), Feinberg (1978), Raz (1986), Thomson (1980; 1990), Kramer (1998), Sreenivasan (2010), Hedahl (2013), and Ripstein (2016). For work more focused on the accountability relation, see Darwall (2006; 2013), Vandieken (2019), Wallace (2019), and Martin (2021).

<sup>5</sup> The word 'jade' refers to two chemically different minerals: see Putnam (1975: 241).

<sup>6</sup> However, even my chosen terms are not immune to ambiguity. There is a sense in which the accountability relation constrains: if an action would warrant your resentment, that should morally constrain me against performing it. However, the sense of 'constraint' I have in mind is more specific: it involves a *deontological* constraint, in the sense elaborated in §1.1. Conversely, the constraint relation implies a form of accountability: even if it is permissible for me to infringe your constraint rights, you have the standing to demand an explanation and (in some cases) compensation for my doing so. But, as I argue in §2, this is different from the more full-throated kind of accountability warranted by the accountability relation. So the words 'accountability' and 'constraint', as used in my parlance, are themselves partly stipulative technical terms.

## 1 | THE CONSTRAINT RELATION AND THE ACCOUNTABILITY RELATION: CHARACTERISTIC FEATURES

The constraint and accountability relations share the same paradigm cases. (This, I suspect, is why most philosophers have not distinguished them). The most-cited case is promising. If I promise you that I will send you comments on a draft of yours by Friday, then other things equal, I am now obligated to you to send you comments by Friday in both the constraint and accountability senses. Promises have a rich suite of normative upshots, all of which have been taken to be characteristic features of directed obligation. If I am right, however, we can divide these features into two clusters, one grounded in the constraint relation, the other in the accountability relation.

A clarification: I understand the features listed below to be *symptoms* of these relations, grounded in and explained by their presence. To delineate these features is not, in my view, to give a theory of what these relations are. Understanding the deeper nature of the constraint and accountability relations is a further project which this paper leaves untouched.

### 1.1 | Features of the constraint relation

The following features of the constraint relation are presented roughly in order of importance, with the most important features first.

(1) *Reason-giving force*. The first thing to observe is that my promise gives me a weighty normative reason to send you comments by Friday. However, as I will emphasize in §2, this reason is overridable. If my child has a medical emergency on Friday and I need to take them to the hospital, then I ought all things considered to break my promise to you. This is why I call the constraint relation an *input concept*: constraint obligations provide normative reasons that are weighed against other reasons in determining what one ought to do.

(2) *Non-instrumentality / action rather than outcome focus*. The constraint relation doesn't just give any sort of reason: my reason to keep my promise has the structure of a deontological constraint. As Shelly Kagan puts it, constraints "erect moral barriers to the promotion of the good" (Kagan, 1998: 73). This means that it can be the case that I ought to keep my promise even if breaking it would produce a better state of affairs. This is why promises provide one of the classic counterexamples to act consequentialism. More generally, my reason to keep my promise is *non-instrumental*: it does not depend on some further valuable outcome that my promise-keeping promotes. This is widely held to be true of the other paradigm constraint rights: I should not kill one to save five; a doctor should not operate on a patient without consent even if doing so will promote their health; I should not read your diary without consent even if you will never find out and be harmed by my doing so.

Another way of describing the same feature is that constraint obligations are *action-focused* rather than *outcome-focused*. Whether one abides by or violates a constraint obligation is a matter of the nature of the action one performs, rather than the (actual or expected) value of the outcome that it brings about. Whether I violate your bodily rights is a matter of whether my action is one of touching or interfering with your body without your consent, not a matter of whether my doing so causes a good or bad outcome.

Notice that only deontologists (broadly understood) hold that there are such non-instrumental, action-focused obligations to perform actions even when they have suboptimal outcomes. As I am characterizing it, the constraint relation is an inherently deontological concept. If act

consequentialism is true, the constraint relation is never instantiated. In contrast, I will argue in §3 that the accountability relation does not presuppose deontology: act consequentialists can hold that some wrongful actions make resentment fitting for particular people.<sup>7</sup>

(3) *Agent-relativity*. The second key deontological feature of the reasons provided by the constraint relation is that they are *agent-relative*. Agent-relative reasons, on a common way of characterizing them, are reasons that can only be described by referring to the agent for whom they are reasons.<sup>8</sup> I have reason to keep *my* promises (agent-relative), not just to see to it that promises are kept (agent-neutral). Deontological constraints are widely understood to provide reasons of the former, agent-relative kind, which is another reason why they resist consequentialist analysis.

While constraints generate agent-relative reasons, not all agent-relative reasons are based in constraints. Some agent-relative reasons are outcome-focused, such as my reason to promote my friend's welfare. This reason is agent-relative because I have partial reason to promote the welfare of *my* friends, even if doing so is not impartially optimal for the welfare of everyone. But it is an instrumental reason to do whatever actions best promote my friend's welfare, not a non-instrumental reason to do or avoid a specific type of action.

(4) *Secondary duties*. Suppose I break my promise to you: Saturday morning comes and I haven't sent you comments. This fact places me under new, 'secondary duties' to you: an obligation to explain myself and to compensate for my breach – say, by sending you extra-thorough comments as soon as possible (see Thomson, 1980 and Ripstein, 2016).

Crucially, I am bound by these secondary duties even if my promise-breaking was permissible all things considered. Again, suppose I didn't send comments because my child fell ill. In this case, I was justified in breaking my promise to you, but I am still under an obligation to explain why I broke my promise and to try to make it up to you. So secondary duties are not merely the result of acting wrongly. Instead, they are the moral residue left behind by the infringement of a constraint obligation – whether that infringement was justified or not.

(5) *Waivability*. Suppose you call me up and say, "I'll have a new draft next week, so don't bother to read this one." By doing this, you've released me from my promise: I am no longer obligated to send you comments by Friday. In other words, you have the power to *waive* my promissory obligation by consenting to my not doing as I promised. This is also a feature of many other constraint rights, such as bodily and property rights. It is normally wrong for me to cut open your skin, but you can make this permissible with consent, as you might if I were your surgeon.

The power to waive is one of the most salient features that distinguishes the person to whom a constraint obligation is owed from everyone else. There is, of course, a dispute about whether this is a feature of all constraint rights, often put as the question of whether there are any 'inalienable rights'.<sup>9</sup> I will argue in §4 that this paper's distinction helps to clarify this debate. But for now we can be neutral on this question and simply observe that waivability is a characteristic feature of many, if not all, constraint obligations. Since it is not obvious that waivability is a universal feature of constraint obligations, however, this should be counted as the least important feature on our list.<sup>10</sup>

<sup>7</sup> Here I am disagreeing with Darwall (2006: 37-38) and Wallace (2019: 86-88).

<sup>8</sup> There are other ways to characterize the distinction, but the differences between them do not matter for our purposes.

<sup>9</sup> See the citations in footnote 34.

<sup>10</sup> Thanks to a reviewer for prompting me to clarify this.

To summarize: when I am under a constraint obligation to you to  $\phi$ , this gives me a non-instrumental, action-focused, agent-relative normative reason to  $\phi$  which can (often but perhaps not always) be waived by your consent, and if I fail to  $\phi$ , I am subject to secondary obligations to explain myself and compensate you, even if my not- $\phi$ ing was permissible all things considered.

## 1.2 | Features of the accountability relation

My promise to send you comments by Friday will also typically place me under an accountability obligation to you. What are the features of this relation?

(6) *Fitting demands and resentment.* My promise to you gives you warrant to hold me accountable to my promise by *demanding* that I fulfill it and *resenting* me if I break it without excuse or justification (Darwall, 2006). If you see me playing video games on Friday, not having started the comments I promised, you can fittingly demand that I get to work. And if I fail to send you comments without excuse or justification, you can fittingly resent me for this.

Let me clarify what I mean when I say that demands and resentment are ‘warranted’ or ‘fitting.’ I will use these two words interchangeably to denote the status a speech act or attitude has when it satisfies its own internal standards of justification. For example, an assertion that  $p$  is warranted only if the speaker has adequate evidence for  $p$ ; it is fitting to fear an object only if it is dangerous. Similarly, demands and resentment are governed by norms that dictate the conditions under which they are warranted or fitting. Crucially, a speech act or attitude can be fitting (/unfitting) even when there are compelling practical reasons to avoid it (/have it) (Rabinowicz & Rønnow-Rasmussen, 2004). A hostage might have warrant to demand that her captor release her even if she ought not make this demand, because doing so will only bring her more abuse. The same point applies to attitudes: resentment can be fitting even if there are strong practical reasons to try to avoid it.

A second clarification regards the distinction between *personal* and *impersonal* forms of accountability. It is sometimes fitting to hold a person accountable for an obligation they owe to someone else. If I start physically attacking you, it is true not only that you can fittingly resent me and demand that I stop, but also that a third party bystander can fittingly feel indignation towards me and demand that I stop. And so the warrant to hold accountable is not always unique to the person to whom an accountability obligation is owed.<sup>11</sup> However, there is an important difference between the personal way in which you can hold me accountable for attacking you, and the impersonal way that a third party can do so. The accountability relation I am after is the relation that grounds the fittingness of personal, not impersonal, accountability.

One way to distinguish these two kinds of accountability is by appeal to P.F. Strawson’s famous distinction between the personal reactive attitude of *resentment* and the impersonal reactive attitude of *indignation*. Resentment is “a reaction to injury or indifference” towards myself, based in my expectation of how I should be treated; indignation is “the vicarious analogue of resentment,” based in a generalized expectation of how anyone should be treated (Strawson, 1967/2008: 15–16). While anyone could fittingly feel indignation at the news of my broken promise, only you are in a

<sup>11</sup> It’s possible that the warrant to hold accountable is *never* unique: that third parties always have warrant to hold me accountable to my obligations to you (for a defense of this claim, see Darwall, 2013; for arguments against it, see Martin, 2021 and de Kenessey, ms). But even if the fittingness of personal and impersonal accountability always go together, they are still distinct and it is still worthwhile to theorize about the relation that makes personal accountability fitting. Thanks to a reviewer for prompting this clarification.

position to fittingly respond with resentment. It is a further question how to unpack this distinction – how is the underlying nature of resentment different from that of indignation? – but all we need here is an intuitive grip on the difference.

It is harder to distinguish between personal and impersonal demands.<sup>12</sup> On the face of it, a demand is the same speech act whether performed by the victim of an attack ('Get your hands off me!') or a bystander ('Get your hands off him!'). We might draw the distinction by appeal to the demands' basis. When you demand that I keep my promise, you are speaking as the person to whom my promise-keeping is owed. When a bystander makes the same demand, she is best understood as making a demand on your behalf, or perhaps on behalf of the moral community (Darwall, 2013: 27). But as I think that personal accountability can be based on agent-neutral reasons (see §3), I don't want to lean on this idea too heavily. So I allow that it may be the case that there is no deep difference between personal and impersonal demands. However, the difference between resentment and indignation, and the uniqueness of the power to forgive (see below), is enough to distinguish personal from impersonal accountability.

Finally, the violation of an accountability obligation is necessary, but not sufficient, for warranted resentment. The connection between the two might be severed if I am not responsible for my wrongful action. In other words, accountability wrongings warrant resentment in the absence of an excuse (Darwall, 2013: 34).

(7) *Power to forgive.* A final important symptom of the accountability relation is the power to forgive. Forgiveness is widely understood as the voluntary relinquishing of resentment (Hieronymi, 2001). As the only person in a position to fittingly resent my promise-breaking, you alone have the power to decide whether and when to give up this resentment. Others may decide to give up their indignation towards me, but this will not make my action count as forgiven. This provides another way of distinguishing between personal and impersonal accountability: even if third parties can hold me accountable for wronging you, only you have the power to forgive me.

One way to make this vivid is to consider apologies. If I am going to apologize for breaking my promise, to whom should I direct this apology? The obvious answer is you. But by directing my apology to you, I implicitly recognize the fact that you alone have the power to accept this apology and forgive me for what I have done (Darwall, 2013: 30–31; Jonker, 2020).

To summarize: when I am under an accountability obligation to you to  $\phi$ , this gives you warrant to demand that I  $\phi$  and resent me if I fail to  $\phi$  without excuse, and the power to decide whether to forgive me for my failure to  $\phi$ .

### 1.3 | How to separate the two relations

Many philosophers have taken all of the features reviewed in this section to be explained by a single, unified cause: my directed obligation to you to keep my promise. My claim is that there are two distinct underlying relations that ground different clusters of these features.

Here's how to test these hypotheses against each other. On the 'one relation hypothesis', we should expect features (1)-(7) to come as a package deal. If these features are all grounded in a single relation of directed obligation, then they should be present when the relation holds, and absent when it doesn't. There might be abnormal cases in which some characteristic features go missing while the rest remain, but such cases should look like isolated exceptions rather than systematic patterns. In contrast, my 'two relation hypothesis' should lead us to expect that the features

<sup>12</sup> Thanks to a reviewer for pointing this out.

listed in this section will come in two dissociable packages (1-5 and 6-7), traceable to two different underlying relations. So, this hypothesis predicts that when we look beyond paradigmatic cases of promissory obligation, we will find examples in which one package of features arises without the other. The next two sections argue that this prediction is correct.

## 2 | CONSTRAINTS WITHOUT ACCOUNTABILITY

This section argues that it is possible to be under a constraint obligation without being under a corresponding accountability obligation. Such cases arise when constraint obligations are *overridden*: when other reasons make it permissible, all things considered, to infringe a constraint obligation.

### 2.1 | The argument

To fix ideas, let's focus on a case of permissible rights violation introduced by Joel Feinberg:

**BLIZZARD:** "Suppose that you are on a backpacking trip in the high mountain country when an unanticipated blizzard strikes the area with such ferocity that your life is imperiled. Fortunately, you stumble onto an unoccupied cabin, locked and boarded up for the winter, clearly somebody else's private property. You smash in a window, enter, and huddle in a corner for three days until the storm abates. During this period you help yourself to your unknown benefactor's food supply and burn his wooden furniture in the fireplace to keep warm. Surely you are justified in doing all these things, and yet you have infringed the clear rights of another person" (Feinberg, 1978: 102).

The two claims in Feinberg's final sentence are key. First, you are justified in breaking into the cabin: all things considered, that is what you ought to do. Second, your breaking into the cabin infringes the owner's constraint rights – specifically, their property right that you not use their cabin without permission.<sup>13</sup> The fact that you have overriding reason to break into the cabin doesn't make the owner's right go away (as I will argue in more detail shortly). So, at least at first glance, it appears to be the case both that you ought to break into the cabin all things considered and that you have a constraint obligation to the owner not to break into the cabin.

Now we can ask: do you have an accountability obligation to the owner not to break into the cabin? I submit that the answer is no, precisely because your breaking into the cabin is justified. In general, one can be under an accountability obligation to  $\phi$  *only if* one ought to  $\phi$  all things considered. To see this, consider the characteristic symptoms of the accountability relation: the warrant to demand and resent, and the power to forgive. Suppose that, after you demand that I  $\phi$  or resent me for failing to  $\phi$ , I show you that I had sufficient reason to do some other action  $\psi$  instead. This is a decisive objection to both your demand and your resentment. If I have sufficient reason to  $\psi$ , then I'm justified in disobeying your demand that I  $\phi$ , and plausibly, a demand cannot be fitting if I can justifiably disobey it. As for resentment, Stephen Darwall puts it well: "it is incoherent to

<sup>13</sup> We are interested in the owner's *moral* property right, which (normally) makes it morally wrong to break in without consent, rather than their legal property right, which makes it illegal to break in.



blame someone for wrongdoing while accepting that he had sufficient reason to act as he did" (2013: 34).<sup>14</sup> To hold on to your resentment while acknowledging that my action was justified would be spiteful and irrational. And if you cannot fittingly resent my action, you do not have the power to forgive me for it either.

To make this more concrete, return to BLIZZARD. Suppose you have the owner on walkie-talkie. You ask their permission to use the cabin, but they refuse, demanding: 'Stay off my property!' Is the owner's demand warranted? Clearly not: they can't fittingly demand that you prioritize their property rights over your life.<sup>15</sup> For the same reason, the owner would not be warranted in resenting you for breaking into their cabin to escape the blizzard. Their resentment implies that you ought to have done otherwise, but that's just not so.

Generalizing from this example, here is my argument:

- (P1) It can be the case both that X has a constraint obligation to Y not to  $\phi$  and that X ought to  $\phi$  all things considered.
- (P2) If X ought to  $\phi$  all things considered, then X cannot be under an accountability obligation not to  $\phi$ .
- (C) Therefore, it can be the case both that X has a constraint obligation to Y not to  $\phi$  and that X does not have an accountability obligation to Y not to  $\phi$ .

I see two ways of resisting this argument, corresponding to the two premises.

First, the objector might deny (P2), arguing that you are under an accountability obligation to the owner not to break into the cabin. They might point out that, even though you are justified in breaking into the cabin, you are still answerable to the owner for violating their rights. At minimum you would owe them an explanation and compensation for the damage you caused their property. Moreover, it would make sense for the owner to be upset that you broke into their cabin, and for you to apologize to them for doing so. Doesn't this show that the owner has warrant to hold you personally accountable for infringing their rights, even though you had sufficient reason to do so?

My reply is to insist that the way you are answerable to the owner for a justified rights infringement is different in kind from the way you would be answerable to them for an unjustified rights violation. The owner's constraint right not to have anyone break into their cabin does imply a secondary right to be compensated if this primary right is infringed. So, under normal conditions, the owner will have warrant to demand that you compensate them for damaging their property and resent you if you do not. What they can't fittingly do is demand that you not break into the cabin in the first place or resent you for doing this.

As for the owner's feeling upset and your apologizing, both of these will take a different form than they would if your action was unjustified. While it is perfectly intelligible for the owner to feel upset that their rights were infringed, their attitude must be compatible with acknowledging that the violation was justifiable. If they were to resent you for infringing their rights, you could legitimately protest that you did what you ought to have done. Similarly, while you can apologize to the owner for violating their rights, your apology should be of the kind that acknowledges that your action was harmful while maintaining that it was justified. You should say, "I am sorry that I *had to* break into your cabin," not "I am sorry *that* I broke into your cabin," for you should

<sup>14</sup> Darwall thinks of resentment as a subtype of blame (as do I), and so his claim applies equally to resentment.

<sup>15</sup> Here I am disagreeing with Margaret Gilbert, who holds that, at least for promises, the promisee has warrant to demand that the promise is kept even if the promisor ought all things considered to break it (Gilbert, 2018: 135).

not regret your choice to do so under the circumstances. So, if we focus our attention on the specific kind of answerability that goes with the accountability relation, it is absent in this case.

## 2.2 | The specificationist reply

The second, more popular reply is to deny (P1): the objector could deny that it is possible to justifiably violate a constraint right. Instead, they could argue that in *BLIZZARD* the owner's property rights are null and void. Rights have implicit exception clauses, they might say: fully stated, the owner's property right is a right that no one break into the cabin *unless they must do so to save their own life*, or perhaps more simply, a right that no one break in *without sufficient justification*. So, the owner's constraint right is not overridden: it simply fails to apply.

This view is called *specificationism*. Specificationism is contrasted with *generalism*, which holds that rights such as the cabin owner's have general content, without exceptions built in for justified breaches.<sup>16</sup> Most specificationists hold that there is an exception clause for every case in which an apparent rights infringement is justified, with the result that rights cannot be permissibly infringed: "The content of a right, on [specificationism], is therefore specified so as not to conflict in the first place with justifiable behaviour" (Oberdiek, 2008: 128). If specificationism is true of constraint rights, then, there are no cases for which (P1) is true, and my argument doesn't get off the ground.

In the course of responding to this objection, I will arrive at a more general conclusion: the debate between specificationists and generalists is a prime example of a literature that could benefit from attention to this paper's distinction. I will argue that each view holds true of a different relation: generalism is true of constraint rights and obligations, while specificationism is true of accountability rights and obligations. The upshot is that these two views, properly understood, are talking past each other. Both should acknowledge the reality of the normative status the other is concerned with; their disagreement only regards which status to call a 'right'.

I will start by criticizing specificationism as a theory of constraint rights. It will help to compare the owner's right in *BLIZZARD* with a constraint right that has an exception built in explicitly. Consider

**INTERRUPTION:** I am busy trying to prepare a lecture I am giving tomorrow, and wish to not be disturbed. However, a good friend is visiting from out of town, and I want to greet him when he arrives. So I ask my wife: "Would you promise not to interrupt me, unless our friend arrives?" She says: "Of course, I promise."

My promissory right that my wife not interrupt me does not apply if our friend arrives: the exception is part of the promise.<sup>17</sup> Our question is: is my wife's interrupting me when our friend arrives relevantly similar to you breaking in to the owner's cabin in order to save your life? If specificationism is true of constraint rights, these cases should be similar: in both, the relevant constraint right does not apply to this context.

<sup>16</sup> Defenders of specificationism include Shafer-Landau (1995), Wellman (1995), Oberdiek (2008), Liberto (2014), and Zylberman (2022). Generalists include Feinberg (1978), Thomson (1980), and Montague (2001).

<sup>17</sup> Generalists should be happy to admit that there can be less explicit exceptions to rights (see Liberto, 2014). The generalist's claim is just that these exceptions do not kick in *every time* I have justification for infringing a right.

I see two major contrasts between INTERRUPTION and BLIZZARD. First, in INTERRUPTION, when our friend arrives, my wife's promise does not give her *any* reason not to interrupt me. There is no conflict between promissory and other reasons, because the promissory reasons don't apply. While in BLIZZARD, you do have normative reason not to break in to the cabin, this reason is just overridden. One way to see this is that if we gradually decrease the harm that would befall you if you didn't break in, breaking in eventually becomes impermissible. If you will get frostbite but not die, perhaps it is still permissible to break in; but if you are just uncomfortable but in no danger, then it is impermissible. Whether breaking in is permissible thus depends on the balance between your reason to break in and your reason not to break in. A specificationist about constraint rights cannot appeal to this balance to explain when breaking in becomes impermissible. More evidence for the presence of a reason is that there is something *regrettable* about the fact that you have to break in, while there is nothing regrettable about my wife's interrupting me so I can greet our friend. The best explanation for this contrast is that the owner's property right remains in force, and thus your reason to respect this right gives you reason not to break in.

The second contrast, often emphasized by generalists, concerns secondary duties.<sup>18</sup> In INTERRUPTION, after my wife interrupts me to greet our friend, she does not owe me any explanation, apology, or compensation. While in BLIZZARD, after I break in, I owe the owner an explanation and compensation for any damage I caused. The best explanation for this appears to be that I infringed the owner's rights, and thus incurred secondary duties of compensation.

Both of these contrasts indicate that the owner's property right, understood as a constraint right, applies in BLIZZARD. If there were an exception, we should expect the case to look like INTERRUPTION: I should have no reason not to break in, and no duty to give compensation.

Specificationists have a reply: they argue that we can explain these contrasts by appeal to the owner's *interests*.<sup>19</sup> In INTERRUPTION, I have no interest in not being interrupted when our friend arrives – I want to greet him. While in BLIZZARD, the cabin owner has an interest in their cabin not being broken into. You thus have a reason not to break in because doing so would harm the owner's interests, and you owe compensation to make up for this harm. Crucially, the owner's interest is not the same as their right: their right does not apply, but their interest remains.

There are two different ways to interpret the role of the owner's interests. The first is direct: your reason not to break into the cabin is that *this specific action* will cause harm to the owner's interests. This reason is thus outcome-focused: it is a reason to avoid bringing about states of affairs that are bad for the owner. The trouble is that your reason not to break in is not outcome-focused in this way. You still have reason not to break in even if your doing so will cause *no* harm to the owner's interests: say, if the door is unlocked and you leave the cabin exactly as you found it. Your reason not to break in remains even if the owner has no desire that you not enter, so you do not harm them even on a desire-fulfillment theory of well-being. Indifference is not the same as consent, and so their property right still applies and gives you reason to respect it. You even have reason not to break in if your doing so will benefit the owner all things considered: say, if they will get a big insurance payout because of the break-in. And in each of these cases, you would still owe the owner at least an explanation and apology, if not compensation. As I argued in §1, your reason to respect the owner's property right is *deontological* in structure, and thus persists even when infringing the right will cause no harm or a net benefit.

<sup>18</sup> The compensation objection is raised by Feinberg (1978: 102) and Thomson (1980). For replies, see Oberdiek (2008: 140-144) and Zylberman (2022: 564-568).

<sup>19</sup> See Oberdiek (2008). Thanks to a reviewer for suggesting this reply, and for pushing me to expand my discussion of specificationism.

This motivates the second interpretation: you have reason not to break in not because that specific action will harm the owner's interests, but because the owner has a *general* interest in people not breaking into their cabin as a rule.<sup>20</sup> We might appeal, in a contractualist or rule-consequentialist style, to the idea that the general adoption of a rule forbidding cabin-breaking would be beneficial for cabin owners. Or we might appeal more directly to a person's interest in generally having autonomy or control over their property.<sup>21</sup> Either way, this proposal takes the owner's general interest to ground a non-instrumental reason that counts against even harmless and beneficial break-ins. It is this deontological reason, grounded in the owner's general interests, that counts against your breaking in even when doing so is harmless and permissible all things considered.

However – and this is critical – this second interpretation *amounts to accepting generalism about constraint rights*. For as I've described them, and as generalists like Feinberg and Thomson understand them, constraint rights *just are* normative reasons with a non-instrumental deontological structure that imply secondary duties when acted against. Grounding these reasons in the owner's general interests is offering an explanation of general constraint rights, not eliminating them. If they acknowledge the presence of deontological reasons not to interfere with property, which lack exception clauses and thus apply even to cases of justified infringements, the specificationist has abandoned specificationism about *constraint* rights.<sup>22</sup>

A closer look at what specificationists say, however, indicates that they might be fine with this result. Several specificationists emphasize that rights, as they understand them, are not a contributory input to practical reasoning, but instead are outputs that depend upon the total balance of reasons. Oberdiek is explicit about this, writing "Under specificationism ... the central role of rights in moral argument is conclusory: fundamentally, one argues towards rights, not from them" (2008: 141); and "rights on [specificationism] ... wait for all the relevant reasons to be accounted for before making an entrance" (135).<sup>23</sup>

All of this is correct about rights *in the accountability sense*. As I have emphasized, the accountability relation is an output concept, not contributing to the balance of reasons but resulting from it. For an action cannot be fittingly demanded, nor can its omission warrant resentment, unless it is what one ought to do all things considered. Thus specificationism is true of accountability rights: if we were to try to articulate the content of an accountability right, we would have to build in exceptions for any action that is permissible all things considered, because such actions cannot be subject to warranted accountability. And so accountability rights cannot be permissibly infringed, as the specificationist holds.

Thus I propose that we should accept specificationism for accountability rights and generalism for constraint rights. Each camp can countenance the existence of two normative phenomena. One is an input normative reason with a deontological structure; the other is an all-things-considered output of one's total reasons. The philosophers on each side of this debate differ simply on which of these phenomena they call 'rights' and 'directed obligations'. Recognizing that there

<sup>20</sup> This is the central idea of the interest theory of rights: see Raz (1986) and Kramer (1998).

<sup>21</sup> Thanks to a reviewer for this proposal.

<sup>22</sup> Can the specificationist insist that these general interests cease to apply when breaking in is permissible all things considered? Yes, but then they will be giving up on the dialectical purpose that motivated their appeal to interests in the first place, which was to explain the persistence of the normative reason against breaking in and secondary duties (these were the contrasts between BLIZZARD and INTERRUPTION). If the general interests do not apply to permissible infringements, they cannot be used to explain these data.

<sup>23</sup> For other specificationists endorsing the same idea, see Shafer-Landau (1995: 215) and Wellman (1995: 282).

are two distinct phenomena here allows us to strike a *détente*: if we interpret generalists as talking about the constraint relation, and specificationists as talking about the accountability relation, the two views are not incompatible after all.

Returning to this section's main argument: if specificationism is not a plausible view of constraint rights, as I have argued, then it cannot provide us with an objection to (P1), the claim that one can permissibly infringe a constraint right. Instead, as I have suggested we interpret it, specificationism amounts to an endorsement of (P2), the claim that one cannot permissibly infringe an *accountability* right. The result is that my argument stands. In BLIZZARD, you have a constraint obligation not to break into the cabin, but have no accountability obligation not to do so. So the constraint and accountability relations are extensionally dissociable: the former can obtain without the latter.

The source of this dissociation is the difference between normative inputs and outputs. Constraint rights provide input reasons that contribute to the overall verdict about what one ought to do. So, they can be overridden by other reasons. When this happens, the constraint right does not disappear: the reasons it grounds are still there, just outweighed. In contrast, accountability rights depend on the total output of one's reasons, because what I can fittingly demand of you or resent you for depends on what you ought to do all things considered. So the fact that one has overriding reason to  $\phi$  *does* preclude one from having an accountability obligation not to  $\phi$ . This is why it is possible to be under a constraint obligation without being under a corresponding accountability obligation.

### 3 | ACCOUNTABILITY WITHOUT CONSTRAINTS

This section defends the converse dissociation: it is possible to have an accountability obligation without being under a corresponding constraint obligation.

To see how this is possible, return to the input/output contrast. The accountability relation is similar to the all-things-considered 'ought' in this respect: both are outputs that result from the overall balance of one's normative reasons.<sup>24</sup> When we consider the all-things-considered 'ought', it is clear that this output can be generated by many different kinds of normative inputs. Some 'oughts' are based in constraints: for example, the fact that I ought not touch your hair without your permission. But few would say that constraint rights are the only considerations that bear on what we ought to do. Plausibly, facts about how my actions will affect my own or others' *welfare* are also relevant to what I ought to do. I ought to watch my wife's favorite movie with her, not because she has a constraint right that I do so, but simply because doing this will make her happy. So, the normative output of 'ought' can be grounded in inputs other than constraints.

Analogously, I will argue that accountability obligations can be grounded in normative reasons other than constraint rights, and so can obtain in their absence. In particular, I claim that the fact that  $\phi$ ing will sufficiently promote or avoid harm to X's welfare can place you under an accountability obligation to X to  $\phi$  even if X has no constraint right that you  $\phi$ .<sup>25</sup>

<sup>24</sup> This is not to say that 'ought' and accountability obligation are the same. There are things I ought to do that no one can fittingly demand of me, such as work on this paper.

<sup>25</sup> This shows that the distinction between constraint and accountability obligations is not merely the more familiar distinction between 'prima facie' or 'pro tanto' and 'all-things-considered' obligation, as some might be suspecting after §2. For if I am right, constraint obligations are not the only 'prima facie' obligations, as there are other inputs that can generate accountability obligations.

Let me emphasize from the outset that many accountability obligations *are* based in constraints. When a constraint obligation is not overridden, it typically grounds an accountability obligation. Other things equal, we owe it to others to respect their bodily, property, privacy, and promissory rights in both the constraint and accountability senses. So I am not claiming that *all* accountability obligations are based in welfare, just that *some* accountability obligations are, and this enables them to exist in the absence of constraints.<sup>26</sup>

Start with an example: my obligation to call my mother on her birthday. If I fail to call my mother on her birthday, I do not violate any of her constraint rights. I do not interfere with her body, property, or privacy. Supposing I have not promised to call her on her birthday, I do not violate any promissory rights either. Yet my failing to call my mother on her birthday still wrongs her in the accountability sense: if I don't call, she can fittingly resent me. Afterwards, I should apologize to my mother and ask for her forgiveness, which she would have the unique power to give. This is a case of accountability obligation without constraint obligation: my mother does not have a constraint right that I call her on her birthday, and yet she still has warrant to hold me personally accountable if I fail to do so without excuse.<sup>27</sup>

Other examples of accountability obligations without constraints include the obligation to throw a life preserver to a struggling swimmer, the obligation to support a friend who has just lost a loved one, and the obligation to tell an acquaintance about the glob of mustard on her chin before she heads into a job interview.<sup>28</sup> If I violate these obligations, I do not thereby violate anyone's constraint rights to body, property, privacy, or any kind of promissory commitment. I simply fail to do something that would benefit the other person substantially at little cost to myself. This fact alone gives the other person warrant to hold me personally accountable: to demand that I throw the life preserver, to resent me for not alerting her to the mustard, to forgive me for not being a good friend in a time of need.

The obvious way to resist this argument is to question my assumption that there are no constraint rights at play in these cases. Why not say that my mother *does* have a constraint right that I call her on her birthday?

We should not categorize this and similar cases as constraint obligations because reasons to promote others' welfare do not have the characteristic features of constraints. I will focus on the two most important features of constraint obligations: non-instrumentality and agent-relativity.

**(1) *Non-instrumentality / action rather than outcome focus.*** Reasons to respect constraint obligations are non-instrumental and action-focused, while reasons to promote welfare are instrumental and outcome-focused. As pointed out in §1, my reason to keep my promise to you or respect your bodily rights does not depend on the beneficial effects of my doing so, and remains in force even if violating these obligations would ultimately benefit you more. If you refuse to consent to a surgical procedure, the fact that it will greatly improve your health does not make it permissible for me to operate on you.<sup>29</sup> The signature feature of constraint obligations is that they persist even when the action they forbid would be harmless or even beneficial, not just agent-neutrally but for the person they wrong.

<sup>26</sup> Thanks to a reviewer for suggesting I clarify this.

<sup>27</sup> See Martin (2021) for a similar line of argument; her example is going to her daughter's recital.

<sup>28</sup> This example is personal: on the first day of a summer school program when I was 16, I had a giant glob of mustard on my chin all day and nobody told me.

<sup>29</sup> Thus my claim that considerations of welfare can ground accountability obligations in the absence of constraint rights does not imply that we should prioritize a person's welfare over her constraint rights when the two conflict (or vice versa).

In contrast, my reasons to promote others' welfare are outcome-focused reasons to bring about a valuable state of affairs. Whether these reasons support an action depends not on the intrinsic nature of the action, but on the expected value of its consequences. If the beneficial consequences disappear, so do the reasons. Suppose my mother hates birthdays and wants to be left alone on hers: then I am no longer obligated to call her. Suppose the struggling swimmer is in a swimming class, and is in no danger of drowning: then I am no longer obligated to throw her a life preserver. If my accountability obligation to  $\phi$  is based in the fact that  $\phi$ ing will promote X's welfare, then if I learn that  $\phi$ ing will not promote X's welfare, the obligation dissipates.

This is why there can be accountability obligations without constraint obligations. One does not need to be an act consequentialist to hold that we have reason to promote good consequences and avoid bad ones. Sometimes I have reason to  $\phi$  just because  $\phi$ ing will bring about an outcome that benefits you or prevent an outcome that costs you. If weighty enough, these outcome-focused reasons can ground accountability obligations. If the benefits of my  $\phi$ ing (or the costs of my not- $\phi$ ing) for you are sufficiently large, and the reasons against my  $\phi$ ing are sufficiently weak, then this fact alone can give you warrant to demand that I  $\phi$  and resent me if I choose not to. If I can save your life by throwing a life preserver, or save your career by pointing out a glob of mustard, then my refusal to help you in these ways plausibly wrongs you in the accountability sense. But my refusal to help does not violate any deontological constraint, as we can see from the fact that it would not be wrong if helping would not benefit you.

(2) *Agent-relativity*. The second signature feature of constraint obligations is that the reasons they ground are agent-relative. I have reason to keep *my* promises and respect your rights *against me*, not to ensure that promises are kept and rights are respected. In contrast, I will argue that accountability obligations can, in some cases, be based in agent-neutral reasons.

Two clarifications are immediately necessary. First, many accountability obligations are based in agent-relative reasons: not just those based in constraints, but some based in welfare. Most nonconsequentialists agree that we have agent-relative reasons to prioritize the welfare of our loved ones. So, my accountability obligation to call my mother on her birthday is plausibly based in my especially strong agent-relative reason to promote her welfare.

Second, a reason can be agent-relative while still being instrumental or outcome-focused. Again, my reason to call my mother on her birthday is agent-relative, but it is also outcome-focused: it depends on the premise that calling my mother will benefit her. This is why I tested for outcome focus by asking whether the obligation persists even if the action is beneficial *for the person to whom it is owed*, rather than merely causing the impersonally best consequences. If an obligation based in agent-relative reasons fails this test, it is outcome-focused and thus not a constraint obligation.

My claim is only that there are *some* cases where accountability obligations are based in agent-neutral, outcome-focused reasons.<sup>30</sup> I submit that the obligation to throw a life preserver to a struggling swimmer is such a case. Supposing the swimmer is a stranger to me, my reason to save them is agent-neutral. Every person has reason to promote the same aim of saving their life; I just happen to be in a position to bring this aim about. Unlike constraints, which give me reasons to do or avoid particular actions myself, my reason to save the struggling swimmer does not give me any reason to ensure that *I* am the one who saves them. If I see a lifeguard approaching the swimmer, and realize that my throwing a life preserver would only interfere with the rescue, then I have no reason to do so (compare Kagan, 1998: 27). My reason to save the swimmer thus lacks two features

<sup>30</sup> In making this claim, I am contradicting both Darwall (2006: 7-9) and Wallace (2019: 38), who argue that accountability obligations are necessarily agent-relative.

widely taken to mark agent-relative reasons: we can characterize the reason without reference to me (it is a reason to save the swimmer), and the reason gives the same aim to every person (to save the swimmer). This provides another dissociation between accountability and constraint obligations, as constraint obligations invariably have these agent-relative features.

One might object that my obligation to save the swimmer still has one signature feature of agent-relativity: I should not violate my duty to rescue the struggling swimmer even if, by doing so, I would foreseeably bring it about that several other people will abide by duties of rescue which they would have otherwise violated.<sup>31</sup> I don't agree with this judgment, but more importantly, it is not forced on us by the concept of accountability obligations. It is perfectly coherent to hold that the swimmer could fittingly resent me if I failed to save them for no good reason, but that they could not fittingly resent me if I failed to save them in order to bring it about that others rescue a greater number of struggling swimmers.

Suppose that there are many struggling swimmers, each of whom will die if they are not rescued. But I am a celebrity, and I know that if I rush into the water to save one swimmer, everyone will be so distracted by watching me perform this rescue that they will wrongfully fail to save the many other struggling swimmers in time. If I hold back and draw no attention to myself, however, I know the many other people on the beach will run in and save many other swimmers – but the one I would otherwise have saved will die. My judgment here is that I ought to hold back: my reason to enable the rescue of many swimmers outweighs the reason to save one myself. This is because my reason to save any of these swimmers is agent-neutral, and so my reason to save a swimmer *myself* is no stronger than my reason to enable someone else to save a swimmer.<sup>32</sup> But even if you disagree with this verdict as a substantive normative claim – say, because you are a skeptic about value aggregation – I think you should agree that it is coherent. More importantly, it is compatible with the claim that if there were only one swimmer to save, and I was the only one in a position to save them, then I would have an accountability obligation, owed to the swimmer, to do so. If this is right, then the claim that I have an accountability obligation to the swimmer (in the one-person case) is compatible with the claim that my reason to save the swimmer has none of the characteristic features of agent-relativity.

This leads to a point I mentioned in §1: act consequentialists can accommodate the existence of accountability obligations based in agent-neutral reasons. *Pace* Darwall (2006: 37–38) and Wallace (2019: 86–88), I see no incompatibility between act consequentialism and the existence of accountability obligations. Perhaps all of our reasons are agent-neutral and outcome-focused, but if we disregard these reasons in a way that especially harms some particular individual, then they have the warrant to hold us personally accountable for this. I am not an act consequentialist, but I see no incoherence in this view. This provides another contrast with constraint obligations, which are incompatible with act consequentialism because they are necessarily non-instrumental and agent-relative.

As I emphasized at the outset, the underlying reason for this section's dissociation goes back to the input/output distinction. The constraint relation is a normative input, which we have characterized by appeal to the deontological content of the normative reasons it entails. Because the accountability relation is a normative output, however, it can in principle be grounded in normative inputs with any kind of content. There is of course a content-based distinction between the

<sup>31</sup> Thanks to a reviewer for raising this objection.

<sup>32</sup> Note that the primary reason for action here is not to prevent more wrongs, but to prevent more deaths. The objector is right that preventing more wrongs shouldn't be of much concern to me, but that's because how many people live or die is much more important than how many wrongs are committed.



normative reasons that can make an action warrant your resentment and those that cannot. But the concept of the accountability relation alone does not obviously tell us what reasons are capable of grounding it: that is a further question for substantive normative theory. I have argued that the set of reasons that can ground accountability obligations includes more than just the deontological reasons implied by constraint obligations: it also includes outcome-focused, welfare-based, and even agent-neutral reasons. When these kinds of reasons are sufficiently strong, they can ground accountability obligations in the absence of the deontological reasons implied by constraints.

One final objection. On R. Jay Wallace's theory of directed obligation, *all* directed obligations, including constraint obligations, are based in the interests of individuals (Wallace, 2019). Wallace writes: "there is nothing that we owe to ... moral persons if their interests do not stand to be affected, in one way or another, by our agency" (2019: 147). Taking inspiration from T. M. Scanlon's contractualism (Scanlon, 1998), Wallace argues that individuals' interests serve as "potential bases for objecting to principles that would permit the agent to act in certain ways" (180), and these objections in turn are the ground of directed obligations (or 'claims', in Wallace's parlance).

Wallace's theory could be seen as blurring the purported line between accountability and constraint obligations. Wallace takes the claims that are generated by contractualist reasoning to have some of the features of constraint obligations: they serve as "presumptive constraints" on deliberation (12), "have an agent-relative character" (38) and are "deontic" (49) in a way that excludes consequentialism (88). But he also takes them to have some of the features of accountability obligations, serving as a basis for fitting resentment (83), demands (84), and the power to forgive (91). Most directly, Wallace argues that directed obligations cannot be permissibly infringed (170-176) – something I have argued is true of accountability but not constraint obligations. Wallace thus might be read as a skeptic about the distinction between accountability and constraint obligations, instead folding features of both into a single conception of directed obligation.<sup>33</sup>

However, I think that Wallace's framework is compatible with accepting this paper's distinction. Wallace recognizes a distinction between the inputs to and outputs of contractualist reasoning. I propose that we identify the claims that Wallace takes to be the outputs of contractualist reasoning with accountability obligations, as they make accountability fitting and cannot be permissibly infringed. Since these outputs are Wallace's target, this amounts to interpreting Wallace as offering a theory of accountability obligations. The question then becomes whether Wallace's view is compatible with recognizing constraint obligations as a distinct input to contractualist reasoning. I think it is.

Wallace endorses something close to the interest-based specificationist view we considered in §2. As we saw there, interests can play two different roles in grounding reasons – or, in the contractualist framework, objections to principles. First, the fact that a specific action will cause direct harm to a person's interests could provide an objection to a principle permitting that action. This objection would ground an instrumental, outcome-focused reason not to perform that action, since whether it applies depends on the consequences of the specific action for the objector's interests. Second, a person might have an interest in actions of that kind not being performed *in general*, which would ground an objection to a principle permitting that action even when it causes no independent harm to the person in a specific case. This objection would provide a non-instrumental, action-focused, agent-relative reason not to perform actions of that kind regardless of their consequences. But this is just what I have been calling a constraint obligation. If Wallacians recognize this second category of input to contractualist reasoning, then, they are thereby admitting the existence of constraint obligations in the sense I intend.

<sup>33</sup> Thanks to a reviewer for suggesting that Wallace's view might offer a challenge to my argument.

Wallace himself appears to recognize the category I am calling constraint obligations, calling them “Hohfeldian claim rights” (2). He argues that the notion of directed obligation he is interested in comes apart from Hohfeldian claim rights in exactly the ways I have argued accountability obligations can come apart from constraint obligations: Hohfeldian rights can be permissibly infringed while directed obligations cannot (175); and not all directed obligations are based in Hohfeldian rights (2). Again, this leads me to think that Wallace’s intended quarry is what I am calling the accountability relation. Wallace does differ from me in holding that accountability obligations are necessarily agent-relative and nonconsequentialist; I explained why I think otherwise earlier in this section. But I care more that other theorists recognize the distinction between accountability and constraint obligations than that they agree with me about all of their features. The other main difference between us is that Wallace insists that we should reserve the words ‘claim’, ‘duty’, ‘wronging’, ‘directed obligation’, etc. for one relation - as I interpret him, the accountability relation. I think doing so invites confusion, as there is an equally well-established tradition of using these terms to refer to the constraint relation. But what matters is that we recognize that there are two distinct normative relations and are clear about which we are investigating at any moment. And Wallace’s theory does not bar us from doing this.

To summarize, I have argued that the constraint and accountability relations are extensionally dissociable in both directions: each can obtain in the other’s absence. As I argued in §2, when a constraint right is permissibly infringed, the constraint relation obtains without the accountability relation. And as I have argued in this section, you can fittingly hold me accountable for acting in a way that significantly harms your interests without sufficient justification even when my action does not violate any constraint rights. In such a case, my reason for action does not have the signature deontological structure of constraint obligations: it is instrumental, outcome-focused, and can in some cases be agent-neutral. Thus the accountability relation can obtain in the absence of the constraint relation.

#### 4 | AN APPLICATION: UNWAIVABLE RIGHTS

This section demonstrates the usefulness of this paper’s distinction by applying it to the question of whether any rights are *unwaivable*. As I mentioned in §1, many paradigmatic rights can be waived by the rightholder’s consent: I can make it permissible for you to touch my hair, use my cabin, read my diary, or not do what you promised by consenting to these actions. The question is whether this feature is merely common or universal: are *all* rights waivable by consent, or are there rights we lack the power to waive?<sup>34</sup>

This question has been a central point of contention in the debate between the *will* and *interest* theories of rights (for a review, see Frydrych, 2018). The will theory says that the distinguishing feature of a right-bearer is their unique power to decide whether to waive or enforce the right. Will theorists thus appear committed to the thesis that all rights are waivable, since they take the power to waive to be constitutive of having a right. In contrast, the interest theory claims that the distinguishing feature of a right-bearer is that the right protects their interests. Interest theorists have tended to deny that all rights are waivable and object to the will theory on this basis (e.g., McCormick, 1977; Kramer, 1998). Though the question of whether all rights are waivable is relevant to the interest vs. will battle, the two are distinct. One could hold that all rights are

<sup>34</sup> Here’s a selection of the literature on this question: McCormick (1977), Feinberg (1978), Kuflik (1986), Kramer (1998: 72–78), McConnell (2000), Sreenivasan (2010: 483), Hedahl (2013: 6).

waivable without thinking that this is constitutive of what it is to have a right, or while thinking that rights are in some way grounded in interests. So I will set aside the debate between the interest and will theories and focus directly on the question of whether there are unwaivable rights.

It will not come as a surprise that I think that how we should answer this question depends on whether we understand ‘rights’ in the accountability or constraint sense. I will argue that the thesis that all rights are waivable is more plausible for constraint rights than it is for accountability rights. I won’t attempt to conclusively show that there are no unwaivable constraint rights. Instead, I will argue that the most salient examples of unwaivable rights are best interpreted as accountability rights, and are less effective as counterexamples to the thesis that all *constraint* rights are waivable. The lesson is that the question of whether all accountability rights are waivable and the question of whether all constraint rights are waivable may well have different answers.

To begin, we need to clarify the thesis that all rights are waivable – what I’ll call the *waivability thesis*.

The first clarification, which I borrow from Feinberg (1978: 114–118), is to distinguish between a right’s being *waived* and its being *relinquished*. In a normal case of waiver, I retain the power to withdraw my consent and bring my rights back into force. After I consent to you cutting my hair, I can at any time withdraw my consent and reinstate my right that you not cut my hair. In some cases, however, we can give up our rights permanently. When I give you my copy of *The Realm of Rights* as a gift, I thereby not only waive my right that you not use the book, but permanently relinquish my power to reinstate that right.

The claim that some rights are ‘inalienable’ is sometimes interpreted as the claim that some rights cannot be relinquished. If we interpret the waivability thesis as denying this, it is obviously false. Of course there are rights one cannot relinquish – the right to refuse sex, for instance. But this is perfectly compatible with the claim that all rights are *waivable*. I might have the power to waive a right even if I lack the power to relinquish it. The interesting, not-obviously-false version of the waivability thesis is the claim that all rights are waivable. I will thus avoid the term ‘inalienable’ and focus on whether there are unwaivable rights.

The second clarification is that, when we assess the waivability thesis, we need to exclude descriptions of actions that presuppose their wrongness or lack of consent. Consider the following cheap argument against the waivability thesis: you can’t waive your right that I respect the rights you haven’t waived; therefore, there are unwaivable rights. This argument misses the mark, because any plausible version of the waivability thesis will agree that I necessarily have a right that you respect my unwaived rights. The interesting question is whether any of my rights against *specific, concrete actions* you might perform – actions like touching my body or entering my cabin – cannot be waived by consent. Thus we should restrict the waivability thesis to rights regarding what I will call *nonmoralized actions*: actions described in a way that is neutral on (a) whether they are morally right or wrong and (b) whether they have been consented to.

This point is both crucial and widely overlooked. Many of the purported examples of unwaivable rights offered in the literature are moralized. One clear example is “the claim-right not to be operated upon without informed consent” (Sreenivasan, 2010: 483). But some examples are more implicitly moralized. One of the most frequently cited unwaivable rights is the right not to be enslaved (MacCormick, 1977: 196; Kuflik, 1986: 75; Sreenivasan, 2010: 483; Hedahl, 2013: 6; the locus classicus is Locke, 1690/1980, Ch. 4). But calling a condition ‘slavery’ presupposes that it is nonconsensual. If one has validly consented and continues to consent to a condition, it no longer

counts as slavery.<sup>35</sup> It is possible to consent to the nonmoralized actions that normally accompany slavery: to working in a cotton field, even to being verbally abused or whipped (consider BDSM practices, in which consensual whipping seems fine). A theorist might insist that one cannot waive one's right not to be whipped or subjected to dehumanizing language. Perhaps they are right: but that's what our debate should be about! The question is whether there are any *nonmoralized* actions involved in slavery to which one cannot effectively consent, rather than whether one cannot consent to 'slavery' as such.

The third and final clarification is that we need to restrict our attention to cases where the rightholder is *competent* to give consent, and her consent is *valid*. It is no counterexample to the waivability thesis that one cannot validly waive one's rights when one is very drunk, or under a coercive threat, or a child. Consent is widely understood to have *validity conditions*: conditions that are necessary for consent to be effective in waiving one's rights (see, e.g., Thomson, 1980: 351). Some of these conditions are situational – not being intoxicated, coerced, or deceived – while some are about the general competence of the rightholder – not being a child or severely mentally disabled. When the validity conditions are not met, one's right cannot be waived, but that does not mean that the right is in principle unwaivable.

Let us define the *normal validity conditions* as the conditions that are required for an uncontroversially waivable right (such as a right that others not cut one's hair) to be successfully waived by consent. Then an interesting version of the waivability thesis will say that all rights can be waived *when the normal validity conditions for consent obtain*, while deniers of this thesis will hold that there are rights that consent cannot waive even when these conditions are met.<sup>36</sup>

With these clarifications in mind, we can formulate the waivability thesis as follows:

**The Waivability Thesis:** For any persons X and Y and nonmoralized action  $\phi$ , if X has a right against Y that Y not  $\phi$ , then if X consents to Y's  $\phi$ ing while the normal validity conditions for consent obtain, then Y can  $\phi$  without wronging X.

We can now put our distinction to work. If talk of 'rights' can be used to refer to either the accountability relation or the constraint relation, then there are two versions of the waivability thesis to consider. Let's start with

**The Accountability Waivability Thesis:** For any persons X and Y and nonmoralized action  $\phi$ , if X has an *accountability* right against Y that Y not  $\phi$ , then if X consents to Y's  $\phi$ ing while the normal validity conditions for consent obtain, then Y can  $\phi$  without wronging X *in the accountability sense*.

The symptoms of wronging in the accountability sense are fitting demands, resentment, and the power to forgive. The question then becomes: can I act in a way you have consented to under normal validity conditions and still thereby warrant your resentment?

I think the answer is yes. Consider some examples:

<sup>35</sup> Even if, as in indentured servitude, one initially consented to the arrangement, it counts as slavery if one's subsequent refusals of consent are disregarded.

<sup>36</sup> Interest theorists have taken the fact that there can be rightholders who are never competent to consent – children, severely mentally disabled people, nonhuman animals – to be a major problem for the will theory (MacCormick, 1977: 198; Kramer, 1998: 69–70). However, my aim here is not to assess the will theory, but to assess the waivability thesis directly. And the most interesting version of the waivability thesis is one restricted to competent consenters.

- **EUTHANASIA:** A patient consents to euthanasia solely because they wish to not be a financial burden to their family; knowing this, the doctor agrees to euthanize them.<sup>37</sup>
- **SURGERY:** A patient consents with full information to a surgical procedure that will put them through sustained extreme pain, where the only purpose of the procedure is a minor aesthetic improvement. The doctor agrees to perform the surgery.
- **BAD EMPLOYMENT:** A business owner employs workers in unsafe conditions for meager wages, with the workers' fully informed consent.

These examples correspond to three commonly cited purportedly unwaivable rights: the right to life (McConnell, 2000), the right to not be tortured (Hedahl, 2013: 6), and the right to fair pay and safe working conditions (Sreenivasan, 2010: 483).

In each of these examples, I submit that (a) the normal conditions for valid consent obtain and (b) the agent acts in a way that makes it fitting for the victim(s) to resent them. If this is true, then there are unwaivable accountability rights.

An objector might deny either of these claims. Start with (b). The objector might argue that by consenting to an action, one gives up the authority to hold the agent accountable for that action. They might point out that it would be *incoherent* for, e.g., the patient to consent to the torturous surgery and simultaneously resent the doctor for agreeing to do it. Suppose this is true. Even if consenting to X's  $\phi$ ing commits me to believing that X's  $\phi$ ing would not warrant my resentment, this is compatible with its being the case that X's  $\phi$ ing *does in fact* warrant my resentment. To see this, imagine that the patient later changes his mind and sees that the doctor should not have accepted his consent. Then I think he can fittingly resent the doctor for having agreed to the surgery. So, while it might be incoherent to consent to and resent the same action simultaneously, it can be fitting to resent an action while knowing that one had validly consented to it when it was performed.

The second, more serious objection targets (a). It might be argued that in each of these cases, the consent is not valid. One might argue that EUTHANASIA and BAD EMPLOYMENT are consent under duress, while the consent in SURGERY may be uninformed. I agree there are ways of filling in the details of these cases in which the consent isn't valid. But the objection requires that *whenever* an apparently-consented-to action warrants resentment, the consent must not have been valid. And I'm not convinced of this. What these examples have in common is that the consenter should not have consented, and the recipient should not have accepted their consent. But inadvisable consent can still be valid. Moreover, not all pressure amounts to invalidating coercion, nor does all ignorance amount to invalidating deception. The patients in EUTHANASIA and SURGERY might be clear-eyed adults, understanding the consequences of their actions, not under any outside pressure, and simply be making a mistake. Even if the conditions for consent are as ideal as they can be, the doctor still should not accept the consent and act upon it, and they warrant the consenter's later resentment by doing so.

This argument is not conclusive, of course, but my aim is just to point out that it is *prima facie* plausible to deny the waivability thesis for accountability rights. In contrast, I think the waivability thesis for constraint rights is more defensible:

**The Constraint Waivability Thesis:** For any persons X and Y and nonmoralized action  $\phi$ , if X has a *constraint right* against Y that Y not  $\phi$ , then if X consents to Y's

<sup>37</sup> If you're worried about the possibility of it being fitting for a dead person to resent, suppose the doctor attempts the procedure and it fails. Could the patient fittingly resent them afterwards?

ping while the normal validity conditions for consent obtain, then Y can  $\phi$  without wronging X *in the constraint sense*.

To deny the constraint waivability thesis, it's not enough to show that some consented-to actions can warrant resentment. Instead, we have to show that the distinctively deontological reason against an action implied by a constraint right can persist even if the action is consented to under normal validity conditions. And this strikes me as less obvious.

Return to my suggested examples of unwaivable accountability rights. The question is whether the *deontological* reason against these actions persists after they are consented to. To test for this, we should ask: is the reason not to perform these consented-to actions fully explained by their expected harmful effects (instrumental/outcome focus), or would it persist even if there were no expected harm, or even benefit, to the rightholder (non-instrumental/action focus)? For each example, I will argue that the relevant reasons are outcome-focused. That is, given that the rightholder has consented under normal validity conditions, the permissibility of the action depends solely on whether it brings about a harmful or beneficial outcome. If true, this would indicate that in our three examples, even when the accountability obligation remains in place, the constraint obligation has been successfully waived – for the deontological, action-focused reason the constraint obligation provides no longer is in force.

**EUTHANASIA:** It is wrong to perform euthanasia in this case because it is not in the patient's interests to do so. The patient's interest in continuing to live is far more important than their family's interest in not being financially burdened. This is why the doctor would warrant the patient's resentment by agreeing to euthanize them. All of this changes if euthanasia is genuinely in the patient's interests. In my view, euthanasia can be permissible when the patient's consent is clearly valid and, crucially, death is clearly in the patient's interests.<sup>38</sup>

This suggests that, once the patient has consented under normal validity conditions, whether euthanasia is permissible depends solely on outcome-focused reasons concerning the patient's interests. If that's right, the action-focused, deontological reason implied by the constraint right against being killed seems to be absent. To see this, contrast a case where the patient does not consent: then whether or not euthanasia is in their interests is irrelevant! Without consent, the constraint right remains in force, making euthanasia wrong regardless of its benefits. But when the patient consents, the normative situation fundamentally changes in a way that indicates that their constraint right has been successfully waived. And as we have seen, this change is compatible with the claim that the patient's accountability right remains in force, as they can still fittingly resent the doctor for agreeing to perform consensual euthanasia if it is not in their interests.

**SURGERY:** To say that there is an unwaivable right 'not to be tortured' commits the mistake of moralizing the action description. To call an action 'torture' implies (if it does not entail) that it is nonconsensual and morally wrong. The interesting question is whether one cannot waive the constraint right to not be subjected to extreme pain – a nonmoralized action. Asked this way, the answer is not obvious. It should be uncontroversial that we can waive our constraint right not to be subjected to mild or moderate pain: we are not wronged by consensual root canals. To hold that there is an unwaivable constraint right not to be subjected to extreme pain, one would have to think that there is some threshold amount of pain which we cannot effectively consent to experience. I doubt this. Having one's leg sawed off without anaesthetic involves torture-level

<sup>38</sup> Three out of the five conditions for permissible euthanasia suggested in the Stanford Encyclopedia of Philosophy article on the topic are focused on ensuring that death is in the patient's interests (Young, 2022).

pain, but if doing so is the only way to save a patient's life (as it could be a couple centuries ago), then it can be permissible with their consent.

The doctor's performing the extremely painful surgery in SURGERY is wrong in the accountability sense not because the patient's consent was invalid, but because the surgery is bad for the patient. The minor aesthetic benefit is simply not worth the extreme pain, even if the patient mistakenly thinks otherwise. If the benefits were large enough to justify the pain – as is plausibly true in the 19<sup>th</sup>-century amputation case – then the patient's consent would make an equally painful surgery permissible in both the constraint and accountability senses. This indicates that, once one has secured another's valid consent, the reason not to subject them to extreme pain is outcome-focused and dependent on the costs and benefits to them of doing so.

Again, this is very different from the normative situation if the patient has not consented. Then it is irrelevant whether the surgery is beneficial or harmful, painful or painless: the constraint right remains in force and makes even a beneficial and painless surgery wrong. This indicates that consent waives the constraint right against experiencing extreme pain, transforming the question of whether the surgery is permissible into a matter of cost/benefit analysis.

**BAD EMPLOYMENT:** Meager wages and unsafe conditions may seem to be necessarily harmful to a worker's interests. But we can imagine cases where the harm is negligible or outweighed. It's easiest with meager wages: suppose that a billionaire wants to work at a local bookstore. It's important to her to be an employee rather than a volunteer, but she doesn't need any money and knows the bookstore's budget is stretched. So she agrees to work for \$5 per hour, under the minimum wage. I think there is nothing wrong with the bookstore owner employing the billionaire under these conditions, because doing so will cause her no harm.

As for unsafe work conditions: it can be permissible to hire someone to work in unsafe conditions if the conditions are not *unnecessarily* unsafe – that is, if the benefits to the worker or others are worth the risks. Think of doctors or journalists working in warzones, or athletes in dangerous sports like free solo climbing. For the warzone doctor or journalist, the risk they consent to undertake is worth it for the help they can give to others. For the free solo climber, the purpose and satisfaction they get from climbing is worth risking death, at least by their lights. If I were to pay someone to do these jobs, I do not think I would be wronging them in either the constraint or accountability senses. But if the risk of harm came to outweigh the benefits, or I was in a position to prevent the harm – as is the case in BAD EMPLOYMENT – then the employment would be wrongful.

Again, contrast the case in which the workers did not consent to laboring under these conditions. Then the question of whether the conditions are beneficial or harmful is irrelevant – forcing someone into labor is wrong regardless of its benefits. This indicates that the workers' consent to working under bad conditions does successfully waive their constraint right against forced labor. Then, the question of whether it is permissible to employ people with meager wages or unsafe conditions becomes a question about outcome-focused, welfare-based reasons, not about deontological constraints.

None of the arguments above are decisive, of course, and there may be other, stronger examples of unwaivable constraint rights. But my aim has not been to settle this question. Instead, my purpose has been to show that what considerations are *relevant* to our assessment of the waivability thesis depends on whether we interpret it as concerning accountability or constraint rights. To determine whether there are unwaivable accountability rights, we need to ask whether it is

possible to act with someone's consent under normal validity conditions and still warrant their resentment. To determine whether there are unwaivable constraint rights, we need to ask whether the deontological, outcome-independent reason against an action normally implied by a person's constraint rights ever persists after they consent under normal validity conditions to that action. These are different questions, and they may well have different answers. The debate about whether there are unwaivable rights will become clearer if we tackle these questions separately. And insofar as the debate between the will and interest theories hinges on the possibility of unwaivable rights, this discussion would also benefit from clarifying which relation these theories are meant to be theories of.

## 5 | CONCLUSION

I have argued that the language of directed obligation – ‘obligation to,’ ‘duty to,’ ‘owed to,’ ‘wronging,’ ‘having a right,’ ‘having a claim,’ ‘entitled to,’ etc. – can be used to describe two distinct normative relations, which I have called the constraint relation and the accountability relation. These two relations have different symptoms: the constraint relation gives rise to normative reasons with a deontological, non-instrumental, agent-relative structure; the accountability relation grounds the fittingness of personal accountability in the form of demands, resentment, and the power to forgive. The relations play different normative roles: the constraint relation is an input, generating a defeasible normative reason to be weighed up against other reasons; the accountability relation is an output, depending on the total balance of one's reasons. These different roles lead to different extensions. When the reason provided by a constraint right is outweighed by other considerations, as in Feinberg's blizzard case, the constraint relation can obtain in the absence of the accountability relation. When an action substantially harms or disregards someone's interests without violating their constraint rights, the accountability relation can obtain in the absence of the constraint relation.

I offer this distinction to my colleagues as a tool to use in further theorizing. I have argued that the question of whether all rights are waivable is clarified by distinguishing between accountability rights and constraint rights. We also saw in §2 that this distinction offers a potential resolution of the debate between specificationism and generalism about rights: specificationists are talking about accountability, generalists are talking about constraints. I believe that other debates about directed obligation and rights would similarly benefit from employing this paper's distinction. If we treat directed obligation as a single thing, we are likely to muddle together constraints and accountability. If we instead separate these two concepts of directed obligation and study each on its own, we will attain a better picture of both.

## ACKNOWLEDGEMENTS

For generous and helpful feedback on drafts of this paper, I would like to thank Tom Hurka, Daniel Muñoz, Kieran Setiya, Jonas Vandieken, and especially the two reviewers for this journal, whose feedback helped me improve the paper substantially.

## ORCID

Brendan de Kenessey  <https://orcid.org/0000-0002-4044-543X>



## REFERENCES

- Attie-Picker, M. (2023). Against the entitlement model of obligation. *Canadian Journal of Philosophy*, 53, 138–155. <https://doi.org/10.1017/can.2023.34>
- Cornell, N. (2015). Wrongs, rights, and third parties. *Philosophy & Public Affairs*, 43, 109–143. <https://doi.org/10.1111/papa.12054>
- Darwall, S. (2006). *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge: Harvard University Press.
- Darwall, S. (2013). Bipolar obligation. In *Morality, Authority, and Law: Essays on Second-Personal Ethics I* (pp. 20–39). Oxford: Oxford University Press.
- Feinberg, J. (1978). Voluntary euthanasia and the inalienable right to life. *Philosophy & Public Affairs*, 7, 93–123.
- Frydrych, D. (2018). The theories of rights debate. *Jurisprudence*, 9, 566–588. <https://doi.org/10.1080/20403313.2018.1451028>
- Gilbert, M. (2018). *Rights and Demands: A Foundational Inquiry*. Oxford: Oxford University Press.
- Hart, H. L. A. (1955). Are there any natural rights? *Philosophical Review*, 64, 175–191. <https://doi.org/10.2307/2182586>
- Hedahl, M. (2013). The significance of a duty's direction: claiming priority rather than prioritizing claims. *Journal of Ethics & Social Philosophy*, 7, 1–28. <https://doi.org/10.26556/jesp.v7i3.75>
- Hieronimi, P. (2001). Articulating an uncompromising forgiveness. *Philosophy & Phenomenological Research*, 62, 529–555. <https://doi.org/10.1111/j.1933-1592.2001.tb00073.x>
- Hohfeld, W. (1913). Some fundamental legal conceptions as applied in judicial reasoning. *The Yale Law Journal*, 23, 16–59. <https://doi.org/10.2307/785533>
- Jonker, J. (2020). Directed duties and moral repair. *Philosophers' Imprint*, 20, 1–32.
- Kagan, S. (1998). *Normative Ethics*. Boulder: Westview Press.
- de Kenessey, B. (ms) [Unpublished manuscript]. Not all moral obligations are obligations to individuals, and not all obligations to individuals are moral obligations.
- Kramer, M. (1998). Rights without trimmings. In M. Kramer, N. E. Simmonds, & H. Steiner (Eds.), *A Debate Over Rights: Philosophical Enquiries* (pp. 7–111). Oxford: Oxford University Press.
- Kuflik, A. (1986). The utilitarian logic of inalienable rights. *Ethics*, 97, 75–87. <https://doi.org/10.1086/292818>
- Liberto, H. (2014). The moral specification of rights: a restricted account. *Law and Philosophy*, 33, 175–206. <https://doi.org/10.1007/s10982-013-9183-4>
- Locke, J. (1690/1980). *Second Treatise of Government*. Hackett Publishing Company.
- MacCormick, D. N. (1977). Rights in legislation. In P.M.S. Hacker & J. Raz, *Law, Morality, and Society: Essays in Honour of H.L.A. Hart*, 189–209. Oxford: Clarendon Press.
- Martin, A. (2021). Personal bonds: directed obligations without rights. *Philosophy & Phenomenological Research*, 102, 65–86. <https://doi.org/10.1111/phpr.12620>
- McConnell, T. (2000). *Inalienable Rights: The Limits of Consent in Medicine and the Law*. Oxford: Oxford University Press.
- Montague, P. (2001). When rights conflict. *Legal Theory*, 7, 257–277. <https://doi.org/10.1017/S1352325201073037>
- Oberdiek, J. (2008). Specifying rights out of necessity. *Oxford Journal of Legal Studies*, 28, 127–146. <https://doi.org/10.1093/ojls/gqm028>
- Putnam, H. (1975). The meaning of 'meaning'. In *Mind, Language, and Reality: Philosophical Papers Vol. 2* (pp. 215–271). Cambridge: Cambridge University Press.
- Rabinowicz, W., & Rønnow-Rasmussen, T. (2004). The strike of the demon: on fitting pro-attitudes and value. *Ethics*, 114, 391–423. <https://doi.org/10.1086/381694>
- Rawls, J. (1955). Two concepts of rules. *Philosophical Review*, 64, 3–32. <https://doi.org/10.2307/2182230>
- Raz, J. (1986). *The Morality of Freedom*. Oxford: Oxford University Press.
- Ripstein, A. (2016). *Private Wrongs*. Cambridge: Harvard University Press.
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge: Harvard University Press.
- Shafer-Landau, R. (1995). Specifying absolute rights. *Arizona Law Review*, 37, 209–224.
- Sreenivasan, G. (2010). Duties and their direction. *Ethics*, 120, 465–494. <https://doi.org/10.1086/652303>
- Strawson, P. F. (1967/2008). Freedom and resentment. In *Freedom and Resentment and Other Essays* (pp. 1–28). New York: Routledge.
- Thomson, J. J. (1980). Rights and compensation. *Noûs*, 14, 3–15. <https://doi.org/10.2307/2214886>

- Thomson, J. J. (1990). *The Realm of Rights*. Cambridge: Harvard University Press.
- Vandieken, J. (2019). Bipolar obligations, recognition respect, and second-personal morality. *The Journal of Ethics*, 23, 291–315. <https://doi.org/10.1007/s10892-019-09297-z>
- Wallace, R. J. (2019). *The Moral Nexus*. Oxford: Oxford University Press.
- Wellman, C. (1995). On conflicts between rights. *Law and Philosophy*, 14, 271–295. <https://doi.org/10.1007/BF01000702>
- Zylberman, A. (2022). Moral rights without balancing. *Philosophical Studies*, 179, 549–569. <https://doi.org/10.1007/s11098-021-01670-9>

**How to cite this article:** de Kenessey, B. (2024). Two concepts of directed obligation. *Philosophy and Phenomenological Research*, 1–26. <https://doi.org/10.1111/phpr.13083>