

Distinguishing Directed Obligation

Brendan de Kenessey

Note to UCLA workshop readers: this is an early draft. One sign of this is that the paper does not engage as much with the prior literature as it should. So please forgive any glaring failures to discuss relevant work on the topic – and please also let me know if you see any such omissions! The paper is also a bit over 30 pages, so if you're pressed for time, feel free to stop at the end of section 3.

A standard project in moral philosophy is to try to sort actions into one of three boxes: morally obligatory, morally permissible, and morally wrong. But this basic trichotomy overlooks an important fact: many of our obligations are owed *to* particular people. If you are drowning and I have a life preserver handy, then I am not merely morally obligated to throw it to you: I am obligated *to you* to do so. If I stomp all over your carefully constructed sand castle in a fit of destructive glee, then I have not merely done wrong: I have wronged *you*.

Following Gilbert (2018), I will call obligations *to* particular persons *directed obligations*, as they are directed *toward* someone.¹ When X is obligated to Y to perform an action ϕ , I'll also say that X's obligation to ϕ is *owed to* Y, and that X's failing to ϕ would *wrong* Y.

Directed obligations form much of the moral scaffolding of our everyday lives: I might spend a typical Tuesday fulfilling my obligations to my students (preparing for class), to my colleagues (going to a department meeting), to my parents (calling them), to my partner (doing the dishes), and to my friend (keeping my promise to join her at pub trivia). And since most morally wrong actions have a victim, most wrongs are also *wrongings*. It takes work to think of moral obligations that *don't* involve obligations to some particular person: perhaps the obligation to give to charity, or to recycle.²

The last twenty years has seen a wave of exciting work on directed obligation.³ Despite this recent interest, however, I think the distinctiveness of directed obligation has not been fully

¹ Other terms include 'bipolar obligation' (Thompson 2004; Darwall 2013) and 'relational obligation' (Wallace 2013).

² See Hare (2012).

³ Examples include Scheffler (2004; 2010), Thompson (2004), Darwall (2013), Owens (2012: Ch. 2), Wallace (2012; 2013; 2019), Căbulea May (2015), and Gilbert (2018).

appreciated. The literature on directed obligation has been dominated by two reductive approaches, one that analyzes directed obligation in terms of rights, and another that sees it as a species of moral obligation. My aim in this paper is to show that both of these reductions fail.

On what I will call the *rights account* of directed obligation, to say that I am obligated to you to ϕ is just another way of saying that you have a *right* against me that I will ϕ . I will argue that, if we interpret “rights” in the way they have been understood in much of moral and political philosophy, this equivalence does not hold (§2). Though many directed obligations are based in rights, there is no necessary connection between the two: your having a right that I ϕ is neither necessary nor sufficient to make it the case that I am obligated to you to ϕ .

On what I will call the *moral-obligation-plus account* of directed obligation, to say that I am obligated to you to ϕ is just to say that I am morally obligated to ϕ , *plus* that you are related to this moral obligation in some special way. On this view, then, directed obligations are simply moral obligations that bear some special relation to a particular person. Against this claim, I will argue that I can be obligated *to you* to ϕ without being *morally* obligated to ϕ (§3). If this is right, then not all directed obligations are moral obligations, and so the former notion cannot be reduced to the latter.

Thus my conclusion will be that there is an important normative relation picked out by “obligation to someone” that is not reducible to the more familiar moral concepts of rights and moral obligation. I can be obligated to you to ϕ in this sense without its being the case either that I am morally obligated to ϕ or that my failing to ϕ would violate your rights.

1. Locating directed obligation

Before we can get started, we need to get an independent grip on the phenomenon I am referring to as directed obligation. The easiest way to do this turns out to be by first looking at the more familiar phenomenon of moral obligation. Though I will ultimately argue that moral obligation and directed obligation are dissociable, there are important parallels between them.

Thus a detour to the topic of moral obligation will put us in a position to better understand our primary target, directed obligation.

I begin with a controversial assumption: there is a difference between saying that I am *morally obligated* to ϕ and saying that I *ought* to ϕ . This assumption is controversial because many philosophers treat the two phrases as equivalent: on their way of speaking, a moral obligation is just something I ought to do.⁴ But there is also a venerable tradition of distinguishing between ‘ought’ and moral obligation.⁵ On this usage, some of the things I ought to do are not morally obligatory: I ought to floss every night, take a stroll when it’s beautiful outside, and cultivate my passion for painting, but it would be odd to say that I am *morally obligated* to do any of these things, or that failing to do them would be *morally wrong*.

In my experience, this distinction is polarizing: it strikes some philosophers as obviously right, and others as totally unintelligible. For those of you in the latter camp, let me explain how I understand the distinction. The sense of ‘ought’ I am distinguishing from moral obligation is the all-things-considered, practical ‘ought’. This is the sense of ‘ought’ that answers the deliberative question, “What to do?” To say that I ought to ϕ is to say that ϕ ing is the correct answer to this question, the thing to be done. In contrast, while to say that ϕ ing is morally obligatory may be *in part* to say that ϕ ing is the thing to be done, it is also to say something more. Following Stephen Darwall (2006) and others, I take the question of whether I am morally obligated to do something to be in part about whether other people can legitimately *hold me accountable* for doing or failing to do it. In particular: if I am morally obligated to ϕ , then this typically gives others standing to hold me accountable by demanding that I ϕ and blaming me if I fail to ϕ without excuse.^{6,7}

⁴ Singer (1972) is an admirably clear example: “I use ‘obligation’ simply as the abstract noun derived from ‘ought,’ so that ‘I have an obligation to’ means no more, and no less, than ‘I ought to’” (233).

⁵ Philosophers who distinguish between ‘obligated’ and ‘ought’ include Hart (1955), Anscombe (1958), Baier (1966), Raz (1977), Brandt (1979), Williams (1985), Scheffler (1992), Gilbert (1993), Wallace (1994), Dreier (2004), Thompson (2004), Darwall (2006), Owens (2008), Portmore (2009), Dorsey (2011), and Harman (2016).

⁶ Philosophers who have drawn this link between obligation and accountability include Mill (1861/2001: 48-49), Ross (1939: 163), Williams (1981: 121), Gibbard (1990: 42), Wallace (1994: 63-64), Skorupski (1999: 142), Gilbert (2004: 90), and Darwall (2006: Ch. 5).

⁷ To be more precise, moral obligation makes demands and blame *fitting* (see Rabinowicz and Rønnow-Rasmussen 2004). A speech act or attitude is fitting in this sense just in case it meets its own internal standard of justification: e.g., an assertion that p is fitting only if the speaker has adequate evidence for p ; it is fitting to admire a person only if

This allows us to give an intelligible interpretation of the claim that I am not always morally obligated to do what I ought. When I say that I ought to floss but am not obligated to do so, what I mean is that flossing is the correct answer to the deliberative question “what to do?” but that no one else has standing to demand that I floss or blame me if I neglect my dental hygiene. Even if you doubt that this view of flossing is true, I think it is hard to deny that it is coherent.

For those who are still skeptical, let me offer another argument for distinguishing ‘ought’ from moral obligation. Most moral theories other than classical act consequentialism accept that in many situations we are morally permitted to do a wide range of actions, including actions that are less than maximally saintly. For example, such a theory might say that it is morally permissible for me to spend Saturday relaxing rather than laboring to do as much good as possible. I think anyone who accepts the existence of such *moral options* will be under strong pressure to distinguish between moral obligation and the all-things-considered ‘ought.’ For if there are moral options, then we will often have to choose between actions that are all morally permissible, e.g.: should I spend Saturday going on a hike or bingeing Netflix? If you think that there can be a correct answer to deliberative questions like this, then you will need a normative notion that draws distinctions within the realm of moral permissibility. ‘Ought’ does this job nicely. If we distinguish ‘ought’ from moral obligation, we can say that I *ought* to go on the hike, though it would be morally permissible for me to spend the day bingeing Netflix instead.

Once we’ve distinguished between ‘ought’ and moral obligation, we can ask about the relation between them. We’ve already seen that ‘ought’ does not entail moral obligation (flossing, hiking). But what about the other way round: if I am morally obligated to ϕ , does that entail that I ought to ϕ ? I think the answer to this question is yes. The reason why has to do with moral obligation’s link with accountability. On the view we’re assuming, I am morally obligated to ϕ if

she is admirable. In this same sense, it is fitting to demand that someone ϕ or blame her for failing to ϕ only if that person is obligated to ϕ . Crucially, an attitude or speech act can be fitting even if there are compelling practical reasons to avoid it. It would be fitting for a slave to demand her freedom from her ‘master’, even if she ought not make this demand all things considered, since doing so would only bring her more abuse. I use all of the following phrases to mean that A is fitting: that one *has standing to* A, that one can *legitimately* A, or that Aing is *warranted*.

and only if others can legitimately demand that I ϕ and blame me for failing to ϕ without excuse.⁸ But others can legitimately demand that I ϕ and blame me for failing to ϕ *only if* I ought to ϕ all things considered. To see this, suppose that, after you demand that I ϕ or blame me for failing to ϕ , I show you that I had sufficient reason to do some other action ψ instead. This seems to be a decisive objection to both your demand and your blame. If I have sufficient reason to ψ , then I'm justified in disobeying your demand that I ϕ , and it is hard to see how a demand could be both warranted and justifiably disobeyed. As for blame, Darwall puts it well: "it is incoherent to blame someone for wrongdoing while accepting that he had sufficient reason to act as he did" (2013: 34). Thus we can argue as follows: (1) I am morally obligated to ϕ only if others are warranted in demanding that I ϕ and blaming me for failing to ϕ ; (2) others are warranted in demanding that I ϕ and blaming me for failing to ϕ only if I ought to ϕ all things considered; therefore (C) I am morally obligated to ϕ only if I ought to ϕ all things considered.⁹

One more point before we turn to directed obligation. Though I've emphasized the link between moral obligation and warranted accountability, I do *not* want to claim that this provides us with a satisfactory theory of moral obligation. I do not endorse a 'fitting response' style account on which *all it is* for an action ϕ to be morally obligatory is for others to have standing to demand that one ϕ and blame one for failing to ϕ . This account gets the order of explanation backwards. If others have the standing to demand that I ϕ or blame me for not ϕ ing, they have this standing *because* I am morally obligated to ϕ , not the other way around (cf. Ewing 1959: 96). I think we should hope for an account of moral obligation that captures how this status *explains* our standing to hold each other accountable with demands and blame.

⁸ For ease of phrasing, I will sometimes omit the "without excuse" qualification, treating it as implicit.

⁹ Dorsey (2011) objects to the view that moral obligation entails 'ought' (which he calls *moral rationalism*) on the basis that it implies that moral reasons always override non-moral reasons in determining what one ought to do. But this only follows if we assume that one is always morally obligated to do what the moral reasons most favor. I reject this assumption. (For objections to it that I find persuasive, see Portmore 2009). Instead, I see the claim that moral obligation entails 'ought' not as an assertion of the overriding force of moral reasons, but as a constraint on how moral reasons generate moral obligations. Moral obligation entails 'ought' not because morality always settles what I ought to do, but because I can only be morally *obligated* to do something if I don't have sufficient non-moral reason to do otherwise. Thus we can hold that moral obligation entails 'ought' without claiming that moral reasons always outweigh non-moral reasons in determining what one ought to do.

Thus I suggest we think of warranted demands and blame as the characteristic *symptoms* of moral obligation. Without purporting to provide an analysis of the concept, moral obligation's conceptual link to accountability can serve as a useful diagnostic tool. Compare: "that which causes smoke" is a lousy analysis of the essence of fire. But looking for smoke is a good way to figure out where there is fire. Similarly, "that which warrants demands and blame" is not, in my view, a satisfactory analysis of moral obligation. But looking for warranted demands and blame is a good way to figure out where there are moral obligations.

Here ends the detour. I have spent all this time on moral obligation because analogues of all the above points apply to directed obligation. As with moral obligation, I think the best way to understand directed obligation is via its link with warranted accountability. What distinguishes directed obligation from moral obligation is the *kind* of accountability these statuses warrant.

To see this, return to the example where I stomp on your sand castle. After you've spent hours carefully shaping the turrets, digging the moat, and lining the walls with pretty pebbles, I leap uninvited onto your sand castle, kicking and stomping like Godzilla in Tokyo. Any bystander, seeing what I am doing, would rightly react with outrage. A bystander might legitimately demand, "Hey! Get off that sand castle!" or simply seethe with warranted indignation at what I have done. But the way in which *you* can hold me accountable is more personal. When you demand that I get off your sand castle, you are not simply holding me to my moral obligations; you are demanding what you are owed. And when you blame me for what I have done, the reactive attitude you feel is not merely impersonal outrage at seeing wrong done, but resentment at having been *wronged*.

This is a dramatization of P. F. Strawson's distinction between the personal reactive attitude of *resentment* and the impersonal reactive attitude of *indignation*. (I take resentment and indignation to be subtypes of blame). Resentment is "a reaction to injury or indifference" towards myself, based in my expectation of how *I* should be treated; while indignation is "the vicarious analogue of resentment", based in a generalized expectation of how *anyone* should be treated (Strawson 2008:

15-16).¹⁰ Building on Strawson, Darwall adds that personal and impersonal accountability invoke different kinds of *authority*: when you demand that I get off your sand castle, you are speaking with *individual* authority as the person whom I am wronging; while when a bystander makes the same demand, she is speaking with *representative* authority, standing in for the wider moral community (Darwall 2013: 27).

What matters for our purposes is just that there is an intuitive difference between the way the victim of a wrong can hold the perpetrator accountable and the way a third party bystander can hold them accountable. Call demands and blame of the former kind “personal” and of the latter kind “impersonal.” Then we can give a diagnostic criterion for directed obligation: X is obligated to Y to ϕ *if and only if* Y has the standing to *personally* demand that X ϕ and *personally* blame X if she fails to ϕ without excuse.

As with moral obligation, I do not take this diagnostic criterion to provide a satisfactory theory of directed obligation. First, the same explanatory point applies: you can hold me personally accountable for destroying your sand castle *because* I am obligated to you not to do so, not the other way around. Second, I doubt that the distinction between personal and impersonal accountability will furnish us with a non-circular way of distinguishing between directed and moral obligation. For I suspect that we won’t be able to say what distinguishes resentment from indignation (or personal from impersonal demands) without invoking the distinction between directed and moral obligation: resentment is a response to *being wronged*, while indignation is a response to *wrongdoing*. In other words, my hunch is that the distinction between directed and moral obligation is explanatorily prior to the distinction between personal and impersonal accountability.

However, the order of discovery need not mirror the order of explanation. If we have an intuitive grip on the difference between personal and impersonal accountability, we can use this to distinguish cases of directed obligation from cases of mere moral obligation. It’s smoke and fire again. “That which warrants personal demands and blame” is a lousy analysis of directed

¹⁰ See also Wallace (1994) and Darwall (2013).

obligation. But consulting our intuitions about who has personal standing to demand and blame is a good way to figure out where there are directed obligations, and to whom they are owed.

To this we can add another way to distinguish the person wronged from third party bystanders: the standing to *forgive*. If I violate a directed obligation to you, then *only you* have the normative power to forgive me (see Darwall 2013: 30-31). No third party bystander has the standing to forgive me for wronging you. They can stop blaming me, sure, but their relinquishing of blame won't have the same normative impact as yours. I won't take a stand on what exactly the normative force of forgiveness is; but whatever it is, it seems clear that only the person wronged has the power to bring it about. Another way to see this is to think about apology: it only makes sense to apologize to the person who was wronged, not to anyone else. By addressing an apology to the person who you wronged, you implicitly recognize their unique standing to decide whether to forgive you. So we can add to our diagnostic criterion: X is obligated to Y to ϕ *if and only if* Y has the standing to *personally* demand that X ϕ , to personally blame X if he fails to ϕ without excuse, and to forgive X for this failure at her discretion.

Now let's turn to the relation between directed obligation and the all-things-considered 'ought.' Here the parallels with moral obligation are clear. First, it can be the case that I ought to ϕ without its being true that I am obligated *to anyone* to ϕ . I ought to floss, but I am not obligated to anyone to do so. Similarly, though I ought to go on a hike, I don't wrong anyone if I stay home and watch TV instead.¹¹

Second, and more interestingly, I submit that directed obligation entails 'ought': X is obligated to Y to ϕ *only if* X ought to ϕ all things considered. My argument for this claim is the same as for the parallel claim about moral obligation. If X is obligated to Y to ϕ , then Y has the standing to demand that X ϕ and blame X for failing to ϕ . But Y can legitimately demand that X

¹¹ Some might insist that I am obligated *to myself* to floss, or to go on the hike. I am frankly not sure what to say about obligations to oneself. But I think the best way to approach the topic is to first focus on cases where the person obligated is not identical to the person to whom the obligation is owed, which I take to be the paradigm cases of directed obligation. Once we have a theory of those cases, we can ask whether it can be extended to explain obligations to oneself. I am presently agnostic on whether such an extension is possible. But if it turns out that our best theory of directed obligation *cannot* make sense of obligations to oneself, I would be inclined to take that as evidence that such obligations do not exist, rather than as reason to reject the theory.

ϕ and blame X for failing to ϕ only if X ought to ϕ all things considered. Suppose it turns out that I had good justification for kicking your sand castle: it was the only way to scare off a lethally poisonous snake that was about to attack you! If you accept that this was sufficient reason for me to kick your sand castle, then I think you cannot coherently blame me for doing so. The same goes for demands: if you think I have sufficient reason not to ϕ , then you cannot coherently demand that I ϕ . Generalizing: I am obligated to you to ϕ only if you have the standing to hold me personally accountable for ϕ ing; you have this standing only if I ought to ϕ all things considered; therefore, I am obligated to you to ϕ only if I ought to ϕ all things considered.

We've arrived at an initial characterization of directed obligation. As with moral obligation, my directed obligations are a proper subset of the actions I ought to do all things considered. In particular, my directed obligations are the actions that some individual person has the standing to hold me personally accountable for performing by demanding that I perform them, blaming me if I fail to, and deciding whether to forgive me for this failure. My obligations to X are the actions for which X can hold me personally accountable in this way; my obligations to Y are the actions for which Y can hold me personally accountable; and so on. This is not a theory of the essence of directed obligation: such a theory would need to explain what it is about certain actions that makes them liable to this personal kind of accountability. But this working characterization gives us all we need to distinguish directed obligation from its conceptual cousins: rights and moral obligation.

2. Distinguishing directed obligation from rights

Here is the most widely accepted view of directed obligation:

THE RIGHTS ACCOUNT: for any persons X and Y and action ϕ , for X to be obligated to Y to ϕ *just is* for Y to have a right against X that X ϕ .¹²

¹² See Bentham (1843: 159), Hohfeld (1913: 31-32), Hart (1955: 179), Anscombe (1990: 152), Thomson (1990: 122), Bond (1996: 199-200), Owens (2012: 46), and Gilbert (2014: 217; 2018). I get most of these references from Cornell (2015: 111-112), who also rejects the rights account, though for different reasons than I offer here.

Many philosophers introduce the rights account simply as a statement of how they will use terminology: they will treat “X is obligated to Y to ϕ ” and “Y has a right that X ϕ ” as equivalent. Understood this way, the rights account is trivially true, like any other specification of the meanings of technical terms. The interesting question is whether, by using terminology in this way, we would be missing an important distinction.

Rights have their own life as a subject of study in moral and political philosophy. If we interpret “having a right” in the way it is understood in this literature, then it refers to a determinate normative property. If we interpret “being owed a directed obligation” in the way set out in section 1, it also refers to a determinate normative property. Then we can ask a substantive, nonterminological question: is the normative property rights theorists refer to as “having a right” identical to the normative property we just described as “being owed a directed obligation”? I will argue that the answer to this question is no.

Let us begin by looking at how rights are understood by those who study them. In the tradition I have in mind, rights are taken to be a family of moral constraints that share a characteristic pattern of features.¹³ The main members of this family are bodily rights, property rights, privacy rights, and the rights created by promises and agreements. These rights share a signature set of properties. First, rights have a *directed* or *relational* structure: the basic phenomenon is one person having a right *against* another. Second, rights can be waived by the *consent* of the right-holder: by giving my consent, I can make it permissible for you to touch my body, use my property, not do what you promised, and so on. And third, rights are *non-instrumental*: I ought to respect your rights not as a means to bringing about some other valuable end, but rather simply because they are your rights. This non-instrumental understanding of rights is often conjoined with the deontological claim that it can be wrong to violate rights even if doing so would bring about the best outcome, including by minimizing the total violation of rights (Nozick 1974: 28-33).

¹³ Some examples of work in this tradition are Hart (1955), Feinberg (1970), Shafer & Frye (1977), Waldron (1981), Thomson (1990), Ripstein (2009), Pallikkathayil (2010), Owens (2011), and Gaus (2012).

Some theorists add to this conception the Kantian idea that rights protect autonomy by giving each person a “sphere of discretionary space: that is, a domain within which she is entitled to control what happens” (Pallikkathayil 2010: 133; see also Ripstein 2009). Your rights delineate a space within which others are (almost always) morally required to let you decide what happens, and so it is (almost always) morally wrong for them to interfere with your choices. For example, you are entitled to decide what happens to your hair: thus you have a right against me (and anyone else) that I not touch, cut, braid, or dye your hair. But you can also exercise your control over your hair by *consenting* to my touching, cutting, braiding or dyeing your hair, thereby making it permissible for me to do so. Thus your hairdresser does not violate your rights, so long as he has your consent.¹⁴ Similarly, property rights give you control over what happens to your property, and privacy rights give you control over who knows certain information about you. Even promises can be brought under this framework: when I promise you that I’ll ϕ , I transfer my right to decide whether I will ϕ to you, now giving *you* the right to control whether I will ϕ by choosing whether to hold me to my promise or release me from it (see Shiffrin 2008 and Owens 2012).

If we read the rights account as employing this (arguably canonical) understanding of rights, and interpret directed obligation in the way set out in §1, then the rights account makes a thoroughly nontrivial claim. Understood this way, the rights account claims that, for any X and Y, Y has the standing to hold X personally accountable by demanding that X ϕ and blaming X for failing to ϕ *if and only if* there is an autonomy-protecting, noninstrumental constraint against X’s ϕ ing that can be waived only by Y’s consent. In other words, the rights account claims that you having a right against me regarding an action, in the sense just articulated, is *necessary and sufficient* for you to have the standing to hold me personally accountable for whether I perform that action.

It’s easy to see why many have thought the rights account is true. Rights and directed obligations usually go hand in hand. If you have a right against me, then you will typically have the standing to hold me personally accountable for respecting that right. The most salient examples

¹⁴ I am borrowing the hairdresser example from Dougherty (2013: 735).

of directed obligations are based in rights: plausibly, you have a right that I not destroy your sand castle, and that explains why you can hold me personally accountable for violating that right. There are more theoretical motivations for adopting the rights account as well. Since both rights and directed obligations have a relational structure – both involve one person owing something *to* another – the rights account simplifies our theory by consolidating two relational phenomena into one. Plus, the rights account offers an attractive answer to the question I posed earlier: what feature of an action explains its liability to personal accountability? The rights account can say: the fact that you have a right against me that I ϕ is what explains your standing to personally demand that I ϕ and blame me for failing to ϕ .

Despite the rights account's prima facie plausibility, I will now argue that it is subject to straightforward counterexamples. Y having a right against X that X ϕ is neither necessary nor sufficient to make it the case that Y has the standing to hold X personally accountable for ϕ ing.

My argument starts from the observation that few rights theorists take rights to be the *only* factor relevant to what we ought to do. Instead, rights-based considerations are typically contrasted with *interest* or *welfare*-based considerations. Rights and welfare have independent, and potentially conflicting, moral import. This is why rights are neither necessary nor sufficient to generate directed obligations: because our obligations to others depend on their *interests* as well as their rights.

Begin with the necessity claim, which says that the *only* actions we can hold others personally accountable for are violations of our rights. This claim is false because considerations of welfare can generate directed obligations without the help of rights. Consider 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 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 I submit that my failing to call my mother on her birthday still wrongs her: if I don't call, she will have standing to feel personal blame towards me. Afterwards, it would make sense for me to apologize to my mother and to ask for her forgiveness, which she would have the unique standing to give. Thus we have a

case of directed obligation without rights: my mother does not have a right that I call her on her birthday, and yet she still has the standing to hold me personally accountable for failing to do so.

Other examples of directed obligations without rights include: the obligation to throw a life preserver to a struggling swimmer, the obligation to tell a friend she has a glob of mustard on her chin before she heads into her job interview, the obligation to hold the elevator door for someone rushing to catch it, the obligation to check in on a friend who has just lost a loved one. If I violate these obligations, I do not thereby interfere in anyone's sphere of discretionary control: I do not violate their 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 no cost to myself. This fact alone gives the other person the standing to hold me personally accountable: to demand that I throw the life preserver or hold the door, to blame me for not telling her about the mustard glob, to forgive me for not supporting him in a time of grief. To generalize: even when Y does not have a right that X ϕ , X may still be obligated to Y to ϕ if ϕ ing would bring a sufficient benefit to Y at a sufficiently low cost to X. If there are *any* such directed obligations, they are counterexamples to the claim that rights are necessary for directed obligation.

Fans of the rights account might be tempted to respond to these examples by expanding the use of "right" to include *any* legitimate claim to be treated in a certain way, including those that are based solely in welfare. So, they would say, my mother *does* have a right that I call her on her birthday, my friend has a right that I alert her to the glob of mustard, and so on. I think this way of using the term "right" is actually common in ordinary usage: it wouldn't be strange for someone to say, "she had a right to know about that glob of mustard." But despite its ordinary language credentials, this way of speaking problematically limits our theoretical vocabulary. It amounts to using "having a right" as simply another term to denote what we have been referring to as "being owed a directed obligation." The result is that we no longer have a term for the narrower phenomenon that has been studied under the name "rights": the consent-sensitive, noninstrumental constraints that are importantly distinct from considerations of welfare. Thus we can preserve the rights account only at the cost of impoverishing our theoretical language.

Let us turn to the sufficiency claim, which says that Y having a right against X that X ϕ is *sufficient* to make it the case that X is obligated to Y to ϕ . This claim is false because rights can be *overridden* by other considerations. Only implausibly absolutist versions of deontology hold that it is *never* morally permissible to violate a right. If violating a right will produce a sufficiently large amount of good or avoid a sufficiently large amount of harm, then it can be permissible to violate that right. It seems wrong to kill one person to save five, but to save five thousand? Five million? Five *trillion* lives? Even if you are an absolutist about killing, it is hard to swallow absolutism about less serious rights violations. It seems clearly permissible for me to break my promise to meet you for lunch in order to take my suddenly severely ill child to the hospital; or for me to break into your remote cabin in order to save myself from perishing in an unexpected blizzard (see Feinberg 1978: 102, and Thomson 1980).

Not only is it morally permissible for me to violate your rights in these cases: I submit that I *ought* all things considered to violate your rights in these cases as well. I *ought* to break my promise to care for my ill child. I *ought* to trespass on your property to save myself from death. Given this, we can argue against the sufficiency claim as follows:

- (1) In some cases, it can be the case both that Y has a right that X not ϕ and that X ought to ϕ all things considered.
- (2) If X ought to ϕ all things considered, then X is not obligated to Y to not ϕ .
- (C) Therefore, in some cases, it can be the case both that Y has a right that X not ϕ and that X is not obligated to Y to not ϕ .

Premise (1) is demonstrated by the cases just considered. Premise (2) is equivalent to the claim, defended in §1, that directed obligation entails ‘ought’: i.e., I can be obligated to you to ϕ only if I ought to ϕ all things considered. The conclusion is a denial of the sufficiency claim: Y having a right that X not ϕ is not sufficient to make it the case that X is obligated to Y to not ϕ .

Let’s apply the argument to the blizzard case (which I get from Feinberg 1978: 102). While hiking in the mountains, I find myself in an unexpected and severe blizzard. Woefully unprepared for the inclement weather, I realize that I will very likely freeze to death if I don’t find shelter. To

my luck, I see your cabin, left empty for the winter. Having no way to ask your permission, I have to decide whether to break into your cabin in order to save myself from nearly certain death.

It seems to me clear that both of the following are true: I ought to break into your cabin all things considered, *and* my breaking into your cabin is a violation of your rights. Now we can ask: am I *obligated to you* not to break into your cabin? No. Given that I ought to break into your cabin, you could not legitimately demand that I not do so or blame me for doing so. Suppose that I have you on walkie-talkie: when I ask your permission to use your cabin, you refuse, and demand: “Stay off my property!” Is your demand warranted? I think not: you don’t have the standing to demand that I prioritize your property rights over my life. For the same reason, it would not be fitting for you to *blame* me when you learn that I broke into your cabin to escape the blizzard. It simply doesn’t make sense to blame someone for an action they were fully justified in doing. Thus I submit that by breaking into your cabin, I violate your rights, but I do not *wrong* you.

In response, a defender of the rights account might point out that, even if I am fully justified in breaking into your cabin, I am still answerable to you for violating your rights. At minimum I would owe you an explanation and compensation for any damage that I caused to your property. Moreover, it would make sense for you to be upset that I broke into your cabin, and for me to apologize to you for doing so. Doesn’t this show that you can hold me personally accountable for violating your rights, even when I had sufficient reason to do so?

My reply is to insist that the way I am answerable to you for a *justified* rights violation is different in kind from the way I would be answerable to you for an *unjustified* rights violation. As both Feinberg (1978) and Thomson (1980) observe, your primary right to not have anyone break into your cabin implies a secondary right to be compensated if this primary right is violated. So, even when I am justified in violating your primary right that I not trespass on your property, I may still be required to respect your secondary right to be compensated for this violation. As for your feeling upset and my apologizing, both of these will take a different form than they would if I had unjustifiably violated your rights. While you are perfectly warranted in feeling upset that your rights were violated, your emotion must be compatible with acknowledging that the violation was

justifiable. If you were to *blame* or *resent* me for violating your rights, I could legitimately protest that I did what I ought to have done. Similarly, while I can apologize for violating your rights, my apology will be of the kind that acknowledges that my action was unfortunate while maintaining that it was justified. I should say, “I am sorry that I *had to* break into your cabin,” not, “I am sorry *that* I broke into your cabin,” for I do not (and should not) regret my choice to do so under the circumstances.

Perhaps the sharpest counterexamples to the sufficiency thesis come from cases in which I justifiably violate your rights for the sake of *your own welfare*. For a somewhat fanciful example, suppose that an assassin is out to kill you. I can get you into protective custody, but only if I find you before the assassin does. I have no idea where you are, but I know that you write down your plans for the day in your private diary every morning. Thus the only way for me to find you before the assassin, thereby saving your life, is for me to read your private diary to learn your location.

It seems clear to me that I ought to violate your privacy rights by reading your diary. More tellingly, I am plausibly *obligated to you* to do so. Suppose that I refuse to read your diary, prioritizing your privacy over your life, and as a result the assassin catches you. If by some miracle you survive the encounter, I think you could legitimately blame me for *not* reading your diary to save your life. “How dare you risk my life for such a small reason!” you might say. Thus I submit that this is an example in which I am obligated *to you* to perform an action that *violates* your rights.

Let us consider one more response on behalf of the rights account. I have argued that, in cases where a right is justifiably overridden, one can violate a right without violating a directed obligation. An objector could respond by denying that any rights are violated in these cases. Instead, she could argue, in the blizzard and assassin cases, the relevant right is null and void. Rights have implicit exception clauses, she could argue: fully stated, your property right is a right that no one break into your cabin *unless they must do so to save their own life*; your privacy right is a right that no one read your diary *unless they must do so to save your life*. The emergency-style cases I have presented trigger these exception clauses, making it the case that an action that would normally

violate your rights does not *in fact* violate them. Thus the rights in these cases are not overridden; they simply cease to apply.

While the objector describes a coherent possibility, I submit that this is just not an accurate description of these cases. In the blizzard and assassin cases, your property and privacy rights remain in full force. This is why there is something *regrettable* about my actions: though I have overriding reason to break into your cabin or read your diary, I also have significant reason *not* to do these things, because they violate your rights. This explains why, if I were able to save my life *without* breaking into your cabin (say, there is a public shelter nearby), or save your life *without* reading your diary (say, I can call someone who knows where you are), I would have strong reason to do so. If your rights were null and void in these cases, then I would have no reason not to break into your cabin or read your diary. We also need to suppose that your rights are in force to explain why you have a right to compensation for my breaking into your cabin, or can legitimately feel upset that I read your diary. These are responses to the fact that your rights *were* violated. If my actions did not violate your rights, you would have no grounds for demanding compensation or feeling upset. The right way to describe these cases is: you have a valid right that I not ϕ , and this gives me strong reason to avoid violating this right. However, my reasons to ϕ happen to be stronger, and so I ought all things considered to ϕ , thereby violating your right that I not ϕ .

Zooming out, I believe there is a single basic difference between rights and directed obligations that explains all of the above counterexamples to the rights account. The difference is this: the concept of a right is a *contributory* or *input* concept, while the concept of directed obligation is a *resultant* or *output* concept. Here's what I mean. Facts about rights serve as inputs to practical deliberation, providing normative reasons that contribute to the overall verdict about what one ought to do. The fact that you have a right that I not break into your cabin gives me a strong reason not to do so. But to reach a verdict about what I ought to do, I must weigh this reason against other normatively relevant considerations, such as the fact that I'll die if I don't find shelter. In contrast, facts about what I am obligated to you to do are a result or output we get by weighing up all of the

relevant normative inputs and seeing what they support in total. For what you can demand of me or blame me for depends on the overall balance of the normative reasons that apply to me.

This helps to explain both why rights and directed obligations usually go together and how they come apart. Rights and directed obligations usually go together because facts about rights provide the input consideration that most commonly generates the output of directed obligation. However, because rights are not the only source of normative reasons, they are not the only source of directed obligations. Input considerations based in welfare can generate directed obligations on their own, which is why rights are not necessary for directed obligation. And because rights provide just one consideration among others, they can in principle be overridden by other normative inputs. Since directed obligations depend on the overall balance of normative reasons, the result is that rights are not sufficient for directed obligation. When your rights are outweighed by other considerations, I will have no obligation to you to respect your rights (blizzard), and may even be obligated to you to *violate* your rights (assassin).

I conclude that there are two distinct normative phenomena that share a directed, relational structure. One relation provides a particularly weighty input to the determination of what one person ought to do by giving them strong reason not to interfere with another person's choices within certain domains. The other relation makes one person personally accountable to another for doing some of the actions she ought to do, giving that other the standing to demand that she perform those actions, blame her if she fails, and forgive her for such failures. These normative relations often show up in the same places, but are conceptually and extensionally dissociable. I have been using "rights" to denote the former relation and "directed obligation" to denote the latter. But I don't care how we label this distinction so long as we recognize its existence. The problem with the rights account is that it doesn't give us the resources to do this: it makes room for only one relational normative concept when in fact there are two.

3. Distinguishing directed obligation from moral obligation

So far I have tied directed obligation very closely to moral obligation. I have characterized both concepts in terms of their connections with accountability, and I have argued that both imply the all-things-considered ‘ought’. These similarities suggest a second reductive account of directed obligation. On this view, there is really only one concept of obligation: moral obligation. To say that I am obligated *to you* to ϕ is just to say that I am *morally* obligated to ϕ , and to add that this moral obligation is related to you in some interesting way. So directed obligations are a subtype of moral obligations: namely, those moral obligations that involve some special relation to a particular individual. I’ll call this view

THE MORAL-OBLIGATION-PLUS ACCOUNT: for any persons X and Y and action ϕ , for X to be obligated to Y to ϕ *just is* for X to be morally obligated to ϕ *plus* for X’s moral obligation to ϕ to bear special relation R to Y.

As its name suggests, the moral-obligation-plus account sees directed obligation as an extra feature added on top of moral obligations. I have deliberately formulated the account to be neutral on what this additional feature is, beyond saying that it must involve some relation to the person to whom the obligation is owed. But to put some flesh on the bones of the moral-obligation-plus account, let’s briefly consider what this relation might be.

An initially appealing proposal is that relation R has to do with *special obligations*. Special obligations are obligations that arise out of special interpersonal relationships, such as those between friends, lovers, family, and role relations such as doctor/patient, student/teacher, lawyer/client and so on. They are usually contrasted with *general* obligations, which apply to one’s interactions with people in general, regardless of one’s relationship to them. Since many paradigm directed obligations are also special obligations, one might propose that the feature that converts a moral obligation into a directed obligation is that it arises out of a special interpersonal relationship. However, this proposal fails for a simple reason: we have directed obligations to strangers. Even if you and I have never met, I am still obligated *to you* to refrain from stomping on your sand castle.

You don't need to be my friend, family, student, client, or anything else to have the standing to hold me personally accountable for how I treat you.

A better candidate for relation R is explanatory: X is obligated to Y to ϕ if and only if X is morally obligated to ϕ *because of* reasons that come from Y. The idea is that I am personally accountable to you for abiding by a particular moral obligation when you are the *basis* or *source* of that obligation. So, I am obligated to you not to destroy your sand castle because the reasons why I am morally obligated not to do so come from facts about *your* interests and rights. Painting in broad strokes, I think this is the most plausible version of the moral-obligation-plus account.

However, my objection to the moral-obligation-plus account does not depend on how we interpret relation R. I get off the bus at an earlier stop: the assumption that directed obligation is a subtype of moral obligation. I will argue that not all directed obligations are moral obligations: I can be obligated to you to ϕ without its being the case that I am morally obligated to ϕ . I call such cases *pure directed obligations*. If pure directed obligations exist, then the moral-obligation-plus account sets out on the wrong path by taking directed obligations to be a subset of moral obligations. Instead, directed obligation is a distinct normative status, importantly similar to moral obligation but not reducible to it.

As with the rights account, there is a danger here of falling into a merely terminological dispute over how to use the labels "moral obligation" and "directed obligation." Why does it matter whether we apply the label "moral obligation" to every directed obligation? The key to answering this question, thereby converting our terminological debate into a substantive inquiry, is to look to these concepts' connections to accountability.

On the account presented in section 1, claims about both moral obligation and directed obligation have substantive entailments regarding who has standing to hold whom accountable. The claim that X is morally obligated to ϕ implies that *any morally competent person* has the standing to hold X accountable by demanding that X ϕ and impersonally blaming X if she fails to ϕ without excuse. In other words, moral obligation grants the standing to hold one accountable *universally* to

every person in the moral community.¹⁵ The claim that X is obligated *to Y* to ϕ implies only that *Y* has the standing to hold X accountable by demanding that X ϕ and personally blaming X if she fails to ϕ without excuse.

With these connections in hand, we can see that the moral-obligation-plus account implies a substantive, and eminently contestable, thesis:

NO PRIVACY: for any persons X and Y and action ϕ , if Y has the standing to hold X personally accountable by demanding that X ϕ and blaming X if she fails to ϕ without excuse, then *everyone* has the standing to hold X impersonally accountable by demanding that X ϕ and blaming X if she fails to ϕ without excuse.

I call this thesis *No Privacy* because it asserts that there is no such thing as *private accountability*: a case in which X is accountable to Y *and only Y* for whether she ϕ s. If directed obligation entails moral obligation, as the moral-obligation-plus account assumes, then accountability is always *public*: if I am accountable to some particular person for performing some action, then I am necessarily also accountable to *everyone* for whether I perform that action. In other words, my obligations to you are never just between us; they are also always everybody else's business.

Thus my argument against the moral-obligation-plus account is based in the claim that accountability *is* sometimes private. I can be personally accountable for you for whether I perform some action without *anyone else* having the standing to demand it of me or blame me for its omission. I believe that such cases of private accountability are not only possible, but pervasive.

Suppose you and I are writing a paper together. There's an upcoming conference that our paper would be a perfect fit for, so we are aiming to get it ready in time for the submission deadline. While you have written your sections promptly, I have procrastinated until suddenly the deadline is a week away and I have yet to begin writing a major section we had assigned to me. It seems clear that you have the standing to demand that I get my act together and write my part of our paper by the deadline, and to resent me if I fail to do so.

¹⁵ A possible exception case is hypocrisy: if some person has unrepentantly violated the same kind of obligation that X has, then he may not have the standing to hold X accountable to this obligation.

Does anyone else have this standing? Suppose our colleague Sally hears about my procrastination on our paper. Can Sally legitimately *demand* that I finish my section in time? Would it be fitting for her to feel *indignation* upon learning that I failed to do so? Or suppose I confide in the barista at my local café about my struggles with our paper. Upon learning the facts, could *the barista* legitimately demand that I finish writing our joint paper, or blame me for failing to do so? When I consult my own intuitions, it seems clear to me that it would be inappropriate – even downright weird – for Sally or the barista to demand that I work harder on a project in which they are completely uninvolved, or to react with blame to my failure to do so. Intuitively, while I am obligated to you to do my share of work on our joint paper, this obligation is between you and me alone. If Sally, the barista, or any other third party were to hold me accountable to this obligation, I think I could rightly protest that it is none of their business.

Of course, a third party might *point out* to me that I am obligated to you to write the section in time, and *advise* me to fulfill this obligation. They might also come to judge my character more negatively upon finding out that I have shirked my obligations to you. But these reactions fall short of the kind of accountability that moral obligation warrants. One way to see this is that advice and character judgment are also appropriate responses to imprudent but morally permissible behaviors, such as my bingeing Netflix rather than going out on a hike. You can appropriately point out to me that I ought to go on the hike, advise me to do so, and form a (mildly) negative judgment about my character upon discovering that I did not take your advice. But it would be clearly out of line for you to demand that I go on a hike or blame me for failing to do so, for I am not morally obligated, or obligated to you, to do this. Similarly, while Sally and the barista could appropriately point out my obligation to you, advise me to fulfill it, and judge me negatively for failing to do so, it would be out of line for them to subject me to demands or blame.

Thus we have an example of *pure directed obligation*: I am obligated to you to finish writing our joint paper by the deadline, but I am not morally obligated to do so. I am obligated to you to write my section because *you* can demand that I do so and personally blame me if I fail. But I am not morally obligated to write our joint paper, for *no one else* has standing to demand that I do so or

blame me for failing to. If I fail to write my section by the deadline, then my failure *wrongs you*, but it is not *morally* wrong.

To broaden our diet of examples, here are some other obligations I propose are pure directed obligations:

- My obligation to my partner to scoop our cats' litter approximately half of the time
- My obligation to my childhood best friend to attend his wedding
- My obligation to my community choir to show up for our rehearsals
- My obligation to my colleague to give comments on a draft she asked me to read

First, I submit that these are genuine directed obligations: the people to whom these obligations are owed can legitimately hold me accountable to them. My partner has standing to demand that I do my share of unpleasant household chores such as scooping cat litter; my best friend could fittingly resent me for skipping his wedding without good reason; my fellow choir members could legitimately demand that I attend our rehearsals; and my colleague could legitimately blame me for being unwilling to comment on her work. Second, I submit that no third parties have the standing to hold me accountable to these obligations. It would be out of place for my neighbor to demand I scoop the cat litter, for my colleague to blame me for skipping my best friend's wedding, for my best friend to blame me for not commenting on my colleague's draft, and for anyone outside of the choir to demand that I attend its rehearsals.

Though I think pure directed obligations are common, this is not to deny that there are also many directed obligations that are accompanied by moral obligation. If a directed obligation is supported by sufficiently strong moral considerations, it will warrant public accountability alongside personal accountability. If I merely neglect household chores, that is between me and my partner alone. But if I were to physically abuse my partner, that horrific behavior would clearly be a moral matter. *Anyone* can demand that I not physically abuse my partner and blame me if I were to do so. Similarly, if I am simply unhelpful to my colleague, then I am accountable to her alone for that. But if I were to plagiarize her work, that awful behavior would plausibly cross the threshold of moral wrongdoing, giving anyone the standing to hold me accountable for what I have

done. Directed obligation and moral obligation often do go hand in hand. However, if I am right that directed obligation is sometimes a solo act, that is enough reason to reject the moral-obligation-plus account.

Note that all of my examples of pure directed obligations are also *special* obligations – obligations that arise out of personal relationships. This suggests a hypothesis about how pure directed obligations arise. It is widely held that personal relationships can generate *partial* reasons for their participants: reasons to treat certain goals, values, or persons as more important than they are from a strictly impartial perspective. From an impartial perspective, my mother’s welfare is no more important than anyone else’s; but plausibly, the fact that she is *my* mother gives me reason to give more weight to her welfare than I give to the welfare of strangers when deciding what to do. From an impartial perspective, it does not matter very much whether we submit our joint paper to the conference: after all, if we don’t submit, the slot we would have occupied will be filled by some other philosopher who will likely benefit just as much from the opportunity. Yet the fact that I have agreed to write this paper with you plausibly gives me reason to assign this goal more value than it has from the point of view of the universe.

It is this divergence between partial and impartial reasons that enables pure directed obligations to arise. Plausibly, morality is impartial. From the moral point of view, every person’s projects, relationships, and welfare matter equally. (At minimum, if some people morally matter more than others, this is not because of the relations they bear to *me*). So moral obligations must be based in impartial considerations.¹⁶ In contrast, directed obligations can be based in partial reasons. If my partial reasons to treat your welfare, our relationship, or our joint projects as specially important give me sufficiently weighty reason to ϕ , then I will be obligated to you to ϕ . This gives us a recipe for generating pure directed obligations: describe a case in which my impartial reasons to ϕ are weak or nonexistent, but my partial reasons to ϕ , based in my relationship

¹⁶ By ‘impartial’ I don’t mean consequentialist. My reason to obey the deontological constraint against killing will count as impartial, so long as that constraint counts equally against my killing each person regardless of their relation to me.

with you, are strong. In such a case, I will be obligated to you to ϕ because of my strong partial reasons to do so, but will not be morally obligated to ϕ , because I have little reason to ϕ from the impartial perspective of morality.

Our examples fit this model nicely. From an impartial perspective, my reasons to finish our paper in time are weak; but my relationship with you, as my colleague and coauthor, gives me strong partial reason to see our joint project through. From an impartial perspective, it does not matter very much whether I attend my choir rehearsals; but my relationship with my fellow choir members gives me strong partial reason to show up and sing.

This is not to say that I have *no* impartial reasons to fulfill my pure directed obligations. From an impartial moral perspective, your frustration at my procrastination and my choir members' disappointment at my absence are both bad things. But adopting a purely impartial point of view, these things are not bad enough to bring moral obligation into the picture. Any view that accepts moral options will have to allow that it can be morally permissible to do things that are *somewhat* bad from the impartial perspective, at least in the sense of not being the impartially best action available (compare Harman 2016). My squandering a Saturday on TV rather than hiking is somewhat bad from an impartial point of view: after all, my welfare matters morally, and I would have been happier if I'd gone for the hike. But my couch-potatoing is not impartially *bad enough* to count as morally wrong. Similarly for our examples of pure directed obligation, even if my violating these obligations is somewhat impartially bad, it is not bad enough to count as morally wrong.

This leads to my final argument for the existence of pure directed obligations. This argument starts from the premise that we can have directed obligations to do *impartially suboptimal* actions. To see this, go back to the coauthored paper case. Suppose that, instead of writing my section of our paper by the deadline, I spend the week volunteering at a homeless shelter. From an impartial point of view, this was probably a better use of my time. Assuming that my volunteering made a difference, the reduction in suffering I brought about by volunteering is almost certainly more valuable from an impartial perspective than the submission of a philosophy paper to a

conference. Yet I submit that I still wronged you by choosing to volunteer rather than write our paper. To choose volunteering on the basis of its impartial optimality would be to ignore my strong partial reason to do my part in our joint project. Intuitively, this partial reason is strong enough to make it the case both that I *ought* to write our paper rather than volunteer and that I am *obligated to you* to do so.

However – and here’s the argument – given that volunteering is impartially better than writing, it is hard to see how I could be *morally* obligated to choose the latter over the former. Moral obligations are based in moral reasons, and moral reasons are impartial. So if volunteering is impartially better than writing, then I must have more moral reason to volunteer than to write. I cannot be morally obligated to do something that is less supported by moral reasons (write) over something that is more supported by moral reasons (volunteering). Therefore, I am not morally obligated to write rather than volunteer. Thus we can conclude that I am under a pure directed obligation: I am obligated to you to write rather than volunteer, but I am not morally obligated to do so.

(This argument needs one refinement. On standard deontological views, I *can* be morally obligated to do an action with impartially suboptimal consequences: for example, I am morally obligated not to kill one innocent person even if doing so will prevent five innocent people from being killed. However, any such obligations to act suboptimally will be based in deontological *constraints*, which prohibit certain actions (killing, rights violations, lying, promise-breaking) even when they would have optimal effects. So we can reformulate the argument’s key claim as: if ϕ ing has impartially better effects than ψ ing, and ϕ ing would not violate a deontological constraint, then one cannot be morally obligated to ψ rather than ϕ . If we stipulate that I did not *promise* you that I would write the paper by the deadline, then our case meets this description: by choosing to volunteer rather than write, I would not violate any deontological constraint. If you suspect that there is an implicit promise here, move to one of the other examples described below.)

So, if you accept that there can be directed obligations to do impartially suboptimal actions (when doing the optimal action would not violate a constraint), then you should also accept that

there are pure directed obligations. And once you allow that we have *reasons* to act partially towards some people, as most non-consequentialist moral theories do, it is hard to deny that we can sometimes be accountable to those people for acting on our partial reasons. But any action based in partial reasons is likely to have an impartially better alternative. The upshot is that if partial reasons can generate directed obligations, then many directed obligations will require impartially suboptimal actions. Instead of fulfilling my obligation to my best friend to attend his wedding, I could do more impartial good by donating the cost of the plane ticket to Doctors without Borders. Instead of going to my community choir's rehearsal, I could do more impartial good by tutoring poor students for free. If I am obligated to my friend or to my choir to do these things rather than their impartially better alternatives, then these are pure directed obligations.

Some readers may still find it hard to swallow sentences like “I am obligated to you to ϕ , but not morally obligated to ϕ ” or “ ψ ing would wrong you, but is not morally wrong.” Let me conclude by offering these readers a terminological salve. The substantive claim I have defended in this section is that there are two distinct kinds of obligation, neither reducible to the other: a kind of obligation that warrants impartial accountability, giving all morally competent persons the standing to demand and blame, and a kind of obligation that warrants personal accountability, giving some particular person the standing to demand and blame. I have been calling the former kind “moral obligation” and the latter “directed obligation,” but we could describe the same distinction with different terminology. An alternative way to go, which some readers may find more palatable, would be to use the term “moral obligation” and its cognates to refer to the broader genus of obligation, encompassing both of the species of obligation I have been at pains to distinguish. We would then need a different term for obligations that warrant impartial accountability: I suggest “impartial obligation.” Then we could say that one is morally obligated to ϕ just in case one has *some* kind of obligation to ϕ , whether directed, impartial, or both. Equivalently, ϕ ing is morally wrong just in case ϕ ing violates either a directed or an impartial obligation; and ϕ ing is morally permissible just in case ϕ ing does not violate either a directed or an impartial obligation. A case of pure directed obligation would then be a case in which I am

obligated to you to ϕ but am not impartially obligated to ϕ . On this use of terminology, pure directed obligations would count as moral obligations, and violating them would be morally wrong – consequences that my target readers might find agreeable.

I am happy with this way of speaking: it is just a different set of labels on the same conceptual map. If we use “moral obligation” this way, the letter of the moral-obligation-plus account turns out to be true: directed obligation is just a subtype of moral obligation. But the view we have reached is importantly different from the moral-obligation-plus account in spirit. The initial idea behind the moral-obligation-plus account was that there was one basic accountability-warranting status, moral obligation, to which we could reduce directed obligation. The view I have been defending is that there are *two* basic accountability-warranting statuses, of which directed obligation is one. If we use “moral obligation” as an umbrella term for both, we are no longer thinking of moral obligation as a more fundamental normative status to which directed obligations are reducible. We are treating directed obligation as its own distinctive normative concept, which is what I have been arguing we should do all along.

4. Toward a positive picture

This paper could have been titled “Groundwork for a Theory of Directed Obligation.” My purpose has been not to provide a theory of directed obligation, but to make space for such a theory by clearing away two common misunderstandings of its subject. In my view, the failure to distinguish directed obligation from both rights and moral obligation has muddied the waters significantly. My aim in this paper has been to rectify this mistake, to defend a map of conceptual space that gives directed obligation its own independent territory.

No matter how decisively a war for territorial independence is won, however, the only way to establish a lasting claim to independence is to build a nation. Similarly, however decisive this paper’s arguments may be, the only way to establish a lasting claim to directed obligation’s conceptual independence would be to build a theory. At risk of abusing my metaphorical license, the intuitive judgments and arguments I offered in sections 2 and 3 are like cannon fire: they can

clear territory, but they cannot hold it. A more lasting way to establish the independence of directed obligation would be to provide a theory of what directed obligation *is*, a theory that gave us a deep enough understanding of directed obligation's nature that we could see for ourselves that it is different from rights and moral obligation. In our metaphor, such a theory would be a wall encircling our new territory – or a constitution for our new nation.

To this end, I'll conclude this paper by outlining the shape I think a positive theory of directed obligation should take. There won't be enough space for me to elaborate or defend this theory here. So, my aim is merely to present a hypothesis for future consideration.

We begin where the last section left off: directed obligation is not reducible to moral obligation, but the two are still closely related. Directed obligation and moral obligation are best seen as two species of the same normative genus.¹⁷

What is this genus? A decent first pass answer is *accountability-warranting statuses*. As I emphasized in section 1, the observation that directed obligation and moral obligation give rise to warranted accountability does not give us a theory of what these statuses are, but it does provide a good starting point for such a theory. We can say: moral obligation is *whatever normative status it is* that makes an action fitting to be demanded, and its omission fitting to be blamed, by any member of the moral community; directed obligation is *whatever normative status it is* that makes an action fitting to be demanded, and its omission fitting to be blamed or forgiven, by a particular individual.

The question for our theory thus becomes: what is the feature of an action that *makes* it a fitting object of personal or impersonal accountability? In other words, what facts about an action *ground* and *explain* others' standing to demand its performance or blame its omission?

I think that the answer to this question lies in the *normative reasons* in support of actions. What makes an action liable to others' demands and blame has something to do with the strength and content of the reasons in that action's favor. Why am I obligated to you to not destroy your sand

¹⁷ Here I am reverting to the use of "moral obligation" to refer to the species of obligation that warrants public accountability, rather than to the whole genus. Those who prefer to reserve "moral obligation" for the genus can substitute "impartial obligation" for "moral obligation" for the remainder of the paper.

castle? Because of the strength of my reasons not to do so, and the fact that these reasons are importantly related to *you*. Why am I not morally obligated to floss? Because the *reasons why* I ought to floss don't have the right kind of content and force to generate a moral obligation. Formulating this idea more precisely, we get:

THE REASONS SCHEMA: for any persons X and Y and action ϕ ,

- (i) For X to be morally obligated to ϕ *just is* for ϕ ing to bear relation R to the subset S_M of X's normative reasons.
- (ii) For X to be obligated to Y to ϕ *just is* for ϕ ing to bear relation R to the subset S_Y of X's normative reasons.

The Reasons Schema is not a theory so much as a recipe for one. To generate our theories of moral and directed obligation, we need to fill in the elements of the schema. We need an account of the sets of normative reasons that ground moral and directed obligations (S_M and S_Y), and an account of the relation an action must bear to these sets of reasons (relation R) to count as a moral or directed obligation. If our theory is to be explanatory, we must fill in these parameters in a way that doesn't implicitly appeal to the concepts we are trying to explain. So, for example, it won't do to say "for X to be obligated to Y to ϕ just is for ϕ ing to be *supported enough* by reasons that *involve Y*." This is true, but unhelpful. To unpack "supported enough" and "involve Y", we'll have to refer back to directed obligation: supported enough *to generate a directed obligation* by reasons that involve Y *in the right way to support a directed obligation*. A more explanatory account would give us a conceptually independent grip on the relevant set of reasons and relation of support that make for a directed (or moral) obligation. Ideally, this would put our theory in a position to provide a non-circular answer to the question: why are all and only actions that bear *this* relation to *that* set of reasons legitimately subject to others' (personal or impersonal) demands and blame?

Here's how I think the Reasons Schema should be filled in. Start with relation R. The relation I propose to fill this role is defined as follows:

GIVING DECISIVE REASON (DEFINITION): A subset S of X 's reasons *gives X decisive reason* to ϕ if and only if the reasons in S alone are sufficient to make it the case that X ought to ϕ all things considered.

Spelling this definition out further, a subset of X 's reasons S is sufficient to make it the case that X ought to ϕ just in case, if we weigh *only* the reasons in S against all of X 's reasons not to ϕ (thus ignoring any non- S reasons in favor of ϕ ing), it will still be the case that X ought to ϕ .

To see this in action, go back to our favorite sand castle example. Consider two reasons that count in favor of my not stomping your sand castle: the fact that stomping your sand castle will make you sad, and the fact that stomping your sand castle will get sand in my shoes. Both of these facts are reasons for me not to stomp, but only the former *gives me decisive reason* not to do so. The fact that stomping your sand castle will make you sad is enough on its own to decisively outweigh all of my reasons in favor of stomping (e.g., that I would enjoy stomping your sand castle), making it the case that I ought not stomp all things considered. But the fact that stomping will get sand in my shoes is plausibly not enough on its own to outweigh the fact that I would enjoy stomping your sand castle. To see this, note that if it were *my own* sand castle that I wished to stomp, the enjoyment I would get from stomping it would easily outweigh the mild annoyance of getting sand in my shoes. So if we consider *only* the fact that stomping will get sand in my shoes, that is not enough to make it the case that I ought not stomp, and so this fact does not give me decisive reason not to stomp.

Thus I propose that I am morally obligated to ϕ just in case the moral-obligation-supporting subset of my reasons gives me decisive reason to ϕ , and I am obligated to you to ϕ just in case the obligation-to-you-supporting subset of my reasons gives me decisive reason to ϕ . Note that this account entails that I can be morally obligated or obligated to you to ϕ *only if* I ought to ϕ all things considered, since the claim that some set of reasons gives me decisive reason to ϕ entails that I ought to ϕ . It should not come as a surprise that I find this to be a salutary result.

The remaining question is: what are the subsets of normative reasons that generate moral and directed obligations? What makes some reasons, and not others, capable of supporting moral

obligations or directed obligations to particular people? Again, I will state my proposed answer to this question without trying to defend it.

I believe that the reasons that generate obligations are the reasons we can appeal to in *joint practical deliberation*, the activity of deciding together what to do. When friends argue about what movie to watch, when colleagues split up work on a project, or when fiancées hash out which relatives to invite to their wedding, they are engaged in joint practical deliberation. When we decide together in this way, we must appeal to a shared body of reasons to justify our joint decisions to each other. My hypothesis is that these are the reasons from which obligations arise.

Motivating this proposal is the idea that there is a deep connection between accountability and *interpersonal justification*. For your demand that I ϕ to be warranted, you must be able to justify that demand *to me*. I am liable to your blame for ψ ing just in case I cannot justify my ψ ing *to you*. This suggests that the reasons that warrant accountability – and thus the reasons that ground obligations – are the reasons we use to justify ourselves to each other. To turn this idea into a theory, we need a story about what the relevant activity of interpersonal justification is. That's where joint practical deliberation comes in: the sense of 'justifiability to others' with which obligation is concerned is justifiability within the activity of deciding together what to do.

This approach suggests a natural way to distinguish between moral and directed obligations. The difference between directed obligations and moral obligations comes down to the different perspectives of joint practical deliberation from which they arise. Directed obligations are based in the particular deliberative perspective shared between the person obligated and the person(s) to whom the obligation is owed. If you and I were to decide together what to do, there is a certain body of reasons we could invoke to justify ourselves to each other. My hypothesis is that these reasons are the basis of our directed obligations to one another. Moral obligations, on the other hand, are based in a perspective of joint practical deliberation that includes *everyone*. In other words, the reasons that generate moral obligations are the reasons that can be taken up within a *universal* perspective of joint deliberation, the reasons we would appeal to if everyone were to decide together what to do.

Putting all of this together, we get:

THE JOINT DELIBERATION THEORY OF DIRECTED OBLIGATION: for any persons X and Y and action ϕ , for X to be obligated to Y to ϕ *just is* for the reasons that apply within the perspective of joint practical deliberation occupied by X and Y to give X decisive reason to ϕ .

I present the joint deliberation theory as a proof of concept, to show that it is possible to offer a positive theory of directed obligation that does not subordinate it to rights or moral obligation. For the joint deliberation theory predicts that directed obligation will diverge from rights and moral obligation in exactly the ways we have seen. Rights and directed obligations can come apart because facts about our rights are plausibly not the only reasons that bear upon our joint deliberation. Moral obligations and directed obligations can come apart because the partial reasons that apply within our two-person perspective of joint deliberation may give me decisive reason to perform an action that is unimportant, or even suboptimal, from the impartial perspective of joint deliberation that includes everyone.

I do not pretend to have given an argument for the joint deliberation theory, or even a satisfactory presentation of it. What I do hope to have done is make space for the question that this theory tries to answer.

References

- Anscombe, G. E. M. 1958. Modern Moral Philosophy. *Philosophy* 33: 1-19.
- Anscombe, G. E. M. 1990. On the Source of the Authority of the State. In *Authority*, edited by Joseph Raz, 142-173. New York: New York University Press.
- Baier, Kurt. 1966. Moral Obligation. *American Philosophical Quarterly* 3: 210-226.
- Bentham, Jeremy. 1943. *The Works of Jeremy Bentham, Vol. 3*, edited by John Bowring. Edinburgh: William Tait.
- Bond, E. J. 1996. *Ethics and Human Well-Being*. Malden: Blackwell.
- Brandt, Richard. 1979. *A Theory of the Good and the Right*. Oxford: Oxford University Press.

- Căbulea May, Simon. 2015. Directed Duties. *Philosophy Compass* 10: 523-532.
- Cornell, Nicolas. 2015. Wrongs, Rights, and Third Parties. *Philosophy & Public Affairs* 43: 109-143.
- Darwall, Stephen. 2006. *The Second-Person Standpoint: Morality, Respect, and Accountability*.
Cambridge: Harvard University Press.
- Darwall, Stephen. 2013. Bipolar Obligation. In *Morality, Authority, and Law: Essays on Second-Personal Ethics I*, 20-39. Oxford: Oxford University Press.
- Dorsey, Dale. 2011. Weak Anti-Rationalism and the Demands of Morality. *Noûs* 46: 1-23.
- Dougherty, Tom. 2013. Sex, Lies, and Consent. *Ethics* 123:717-744.
- Dreier, James. 2004. Why Ethical Satisficing Makes Sense and Rational Satisficing Doesn't. In *Satisficing and Maximizing: Moral Theorists on Practical Reason*, edited by Michael Byron, 131-154.
Cambridge: Cambridge University Press.
- Ewing, A. C. 1959. *Second Thoughts in Moral Philosophy*. New York: Routledge.
- Feinberg, Joel. 1970. The Nature and Value of Rights. *The Journal of Value Inquiry* 4: 243-257.
- Feinberg, Joel. 1978. Voluntary Euthanasia and the Inalienable Right to Life. *Philosophy & Public Affairs* 7: 93-123.
- Gaus, G. F. 2009. Recognized Rights as Devices of Public Reason. *Philosophical Perspectives* 23: 111-136.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge: Harvard University Press.
- Gilbert, Margaret. 1993. Agreements, Coercion, and Obligation. *Ethics* 103: 679-706.
- Gilbert, Margaret. 2004. Scanlon on Promissory Obligation: The Problem of Promisees' Rights. *Journal of Philosophy* 101: 83-109.
- Gilbert, Margaret. 2014. The Nature of Agreements: A Solution to Some Puzzles about Claim-Rights and Joint Intention. In *Rational and Social Agency: The Philosophy of Michael Bratman*, edited by Manuel Vargas and Gideon Yaffe, 215-256. Oxford: Oxford University Press.
- Gilbert, Margaret. 2018. *Rights and Demands: A Foundational Inquiry*. Oxford: Oxford University Press.

- Hare, Caspar. 2012. Obligations to Merely Statistical People. *Journal of Philosophy* 109: 378-390
- Harman, Elizabeth. 2016. Morally Permissible Moral Mistakes. *Ethics* 126: 366-393.
- Hart, H. L. A. 1955. Are There Any Natural Rights? *Philosophical Review* 64: 175-191.
- Hohfeld, Wesley. 1913. Some Fundamental Legal Conceptions as Applied in Judicial Reasoning. *The Yale Law Journal* 23: 16-59.
- Mill, John Stuart. 1861/2001. *Utilitarianism* (Second Edition), edited by George Sher. Indianapolis: Hackett Publishing Company.
- Nozick, Robert. 1974. *Anarchy, State, and Utopia*. New York: Basic Books.
- Owens, David. 2008. Rationalism about Obligation. *European Journal of Philosophy* 16: 403-431.
- Owens, David. 2011. The Possibility of Consent. *Ratio* 24: 402-421.
- Owens, David. 2012. *Shaping the Normative Landscape*. Oxford: Oxford University Press.
- Pallikkathayil, Japa. 2010. Deriving Morality from Politics: Rethinking the Formula of Humanity. *Ethics* 121: 116-147.
- Portmore, Douglas. 2009. Are Moral Reasons Morally Overriding? *Ethical Theory & Moral Practice* 11: 369-388.
- Rabinowicz, Wlodek, and Toni Rønnow-Rasmussen. 2004. The Strike of the Demon: On Fitting Pro-attitudes and Value. *Ethics* 114(3): 391-423.
- Raz, Joseph. 1977. Promises and Obligations. In *Law, Morality, and Society: Essays in Honour of H. L. A. Hart*, edited by P. M. S. Hacker and Joseph Raz, 210-228. Oxford: Clarendon Press.
- Ripstein, Arthur. 2009. *Force and Freedom: Kant's Legal and Political Philosophy*. Cambridge: Harvard University Press.
- Ross, W. David. 1939. *Foundations of Ethics*. Oxford: Oxford University Press.
- Scheffler, Samuel. 1992. *Human Morality*. Oxford: Oxford University Press.
- Scheffler, Samuel. 2004. Projects, Relationships, and Reasons. In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, edited by R. Jay Wallace, Philip Pettit, Samuel Scheffler & Michael Smith, 247-269. Oxford: Oxford University Press.
- Scheffler, Samuel. 2010. Morality and Reasonable Partiality. In *Partiality and Impartiality: Morality,*

- Special Relationships, and the Wider World*, edited by Brian Feltham and John Cottingham, 98-130. Oxford: Oxford University Press.
- Shafer, Carolyn M., and Marilyn Frye. 1977. Rape and Respect. In *Feminism and Philosophy*, edited by Mary Vetterling-Braggin, Frederick A. Elliston, and Jane English, 336-346. Totowa: Rowman & Littlefield.
- Shiffrin, Seana Valentine. 2008. Promising, Intimate Relationships, and Conventionalism. *Philosophical Review* 117:481-524.
- Singer, Peter. 1972. Famine, Affluence, and Morality. *Philosophy and Public Affairs* 1: 229-243.
- Skorupski, John. 1999. *Ethical Explorations*. Oxford: Oxford University Press.
- Strawson, P. F. 1967/2008. Freedom and Resentment. In *Freedom and Resentment and Other Essays*, 1-28. New York: Routledge.
- Thompson, Michael. 2004. What Is It to Wrong Someone? A Puzzle about Justice. In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, edited by R. Jay Wallace, Philip Pettit, Samuel Scheffler & Michael Smith, 333-384. Oxford: Oxford University Press.
- Thomson, Judith Jarvis. 1990. *The Realm of Rights*. Cambridge: Harvard University Press.
- Thomson, Judith Jarvis. 1980. Rights and Compensation. *Noûs* 14: 3-15.
- Waldron, Jeremy. 1981. A Right to Do Wrong. *Ethics* 92: 21-39.
- Wallace, R. Jay. 1994. *Responsibility and the Moral Sentiments*. Cambridge: Harvard University Press.
- Wallace, R. Jay. 2012. Duties of Love. *Proceedings of the Aristotelian Society Supplementary Volume* 86: 175-198.
- Wallace, R. Jay. 2013. The Deontic Structure of Morality. In *Thinking About Reasons: Themes from the Philosophy of Jonathan Dancy*, edited by David Bakhurst, Brad Hooker, and Margaret Olivia Little, 137-167. Oxford: Oxford University Press.
- Wallace, R. Jay. 2019. *The Moral Nexus*. Princeton: Princeton University Press.
- Williams, Bernard. 1981. Ought and Moral Obligation. In *Moral Luck: Philosophical Papers 1973-1980*, 114-123. Cambridge: Cambridge University Press.
- Williams, Bernard. 1985. *Ethics and the Limits of Philosophy*. Cambridge: Harvard University Press.