

Obligations of Feeling

1

One influential conception of moral obligation is given by John Stuart Mill. In a famous passage, Mill tells us that

We do not call anything wrong unless we mean to imply that a person ought to be punished in some way or other for doing it—if not by law, by the opinion of his fellow creatures; if not by opinion, by the reproaches of his own conscience...There are other things, on the contrary, which we wish that people should do...but yet admit that they are not bound to do; it is not a moral obligation; we do not blame them (Mill, 2001: 48-9).

The idea is that blame, guilt, and related practices of censure *track* moral obligation and separate it out from other normative concepts. As Mill puts it, “This seems the real turning point between morality and simple expediency” (49).

But there is another crucial bit. Mill writes,

It is a part of the notion of duty in every one of its forms that a person may rightfully be compelled to fulfill it. Duty is a thing which may be *exacted* from a person, as one exacts a debt (49, emphasis in the original).

The claim here is that obligations are enforceable, they can be “exacted from a person, as one exacts a debt.” This last point has fallen out of favor, but a related idea remains popular, namely, that obligations, especially directed obligations, are linked with

demands. As R. Jay Wallace puts it, “the very things that give us reasons to do (say) *X* also give another party a claim against us that we should do *X*” (Wallace: 2012, 194). The obligee, in other words, can demand that we do *X*.

The puzzle is the following: when it comes to our attitudes, in particular to our emotional responses, Mill’s two components seem to come apart. On the one hand, we don’t make demands on others’ emotions. If I promise to do *X*, then you can demand that I do *X*. You can say: “Keep your promise and do *X*!” But to demand others to feel something seems misguided, unintelligible. While we might say things like, ‘You should feel happy about him’ we never say, ‘Feel happy about him!’ Such a remark sounds absurd.

And yet, we do blame others (and ourselves) for their emotions. And we blame them not in the way we blame a natural disaster or a ten-year old for her hurtful words, but much in the way we blame fully responsible adults for their actions. We not only lament that our spouse is jealous, but we *blame* him for it. We don’t simply see our lack of gratitude as a character flaw (something bad, perhaps regrettable, about us), but often think of it as wrong—we *should* feel gratitude, we think, and feel guilt in response.

My claim is, using Gary Watson’s (1996) distinction, that we are responsible for our emotions in the accountability sense (and not simply in the attributability sense). That is, our emotions not only warrant judgments about our character—about who we are as a person—but also the sort of responses philosophers usually associate with blame, guilt, and the rest of the reactive attitudes. Or rather, the claim is that our *actual practices* treat emotions in this way. As a matter of fact, that is, we hold each

other accountable for what we feel. This is evident not only in the fact that blame and guilt (and the rest of the reactive attitudes) are often directed at them. It is also suggested by the related practices of excuse and justification, and the give-and-take of apology and forgiveness. Or so I shall argue.

The following, I claim, is generally true of emotions:

- To demand them is infelicitous.

And yet,

- They are the object of blame and guilt.
- Justifications and excuses are given and expected.
- Apologies are offered and forgiveness is granted (or refused).

The challenge is to sort out the moral status of emotions given these facts. Do we have obligations of feeling, even if we can't demand that they be fulfilled? Or is it rather that the reason we can't demand them is precisely that no one is obligated to feel anything, even though we hold each other accountable for doing so? Either way there is some explaining to do.

One natural response, however, is to claim that we don't hold each other accountable for our emotions, but only for their expression. We don't blame others for how they feel but only for how, and insofar as, they act on those feelings. So no obligations of feeling exist. The puzzle is a pseudo-puzzle, a nonproblem.

Though appealing at first, this response fails to take seriously, in P.F. Strawson's words, "the very great importance that we attach to the attitudes and intentions of other human beings, and the great extent to which our personal feelings and reactions depend upon, or involve, our beliefs about these attitudes and intentions" (Strawson, 1982: 62. Cf. Adams (1985) and Smith (2005: 249-50)).

Consider a real-world example from a few years back:

It was not a pretty sight Sunday, when fans of the Kansas City Chiefs cheered as their team's quarterback, Matt Cassel, had to leave the game against the Baltimore Ravens because of head injury...Cassel remained prone on the ground after being hit, and fans cheered the sight of the injured Cassel, who has not played that well for the Chiefs the last two seasons (Mitchell, 2012).

Cassel's teammate, offensive tackle Eric Winston, said the following after the game:

When you cheer, when you cheer somebody getting knocked down, I don't care who it is...it's sickening. It's 100 percent sickening...If you are one of those people, one of those people who were out there cheering or even smiled when he got knocked out, I just want to let you know, and I want everybody to know that I think it's sickening and disgusting. We are not gladiators and this is not the Roman Coliseum...it's sickening and I was embarrassed. And I want every one of you people to put that on your station, to put that on your newspapers, because I want every fan to know that (cited in Strachan, 2012).

What "sickens" Winston about the fans' response? Is it the cheering, the expressive act? Or rather the emotion or attitude the cheering expresses? In other words, is

Winston blaming¹ the fans for *expressing* their happiness at Cassel's injury² (as the objector suggests)? Or is he also blaming them for *being happy*? It seems to me that at least part of Winston's blame has as its object the fans' emotion, in itself, irrespective of its expression. I think that's exactly what he means when he singles out those "who were out there cheering *or even smiled* when he got knocked out."

Of course, smiling is an expressive gesture, so he could have meant that any expression of happiness, no matter its form, was reprehensible under the circumstances. Note on this score that the cheering is not merely a window to the fans' happiness—a mirror that reflects the emotion outwards into public space. The cheering also expresses an intention to deride the player. It expresses contempt, ill-will. And sometimes it is indeed these attitudes that are the object of blame. Sometimes it is the intention to communicate a given attitude, in itself perhaps unobjectionable, that becomes the subject of censure. Suppose someone tells you that your baby is ugly. What might bother you in this case is not her thought that your baby is ugly (after all, he may be quite ugly), but rather that she had the guts to tell

¹ Some readers have suggested that Winston's language makes it unclear whether he is actually blaming the fans, given that being sickened and disgusted are arguably not reactive attitudes. I encourage those with this reaction to watch a clip of the interview and decide for themselves:

<https://www.youtube.com/watch?v=3mijDWa3CLM&t=76s>. To me, Winston's interview is a *paradigm* case of blaming and holding accountable.

² What's the object of the fans' reaction? There are many possibilities (and, of course, some might be true of some fans and not of others). One is that the fans were simply excited to get a new quarterback in the game. If so, it is a mistake to see their reaction as directed at Cassel's *injury*. Or perhaps the fans were happy for his injury *but only because* that meant a new quarterback was entering the game. Here, the injury is indeed part of what the emotion is about, but the following would still be true: the fans would have been as happy had the coach simply made a substitution and Cassel doesn't get knocked down. A more troubling possibility is to see the fans as enjoying Cassel's injury for its own sake. We can imagine many fans seeing the injury as payback for Cassel's performance ("He had it coming!")—a kind of divine punishment for all those interceptions. A different (though not incompatible) possibility is to see the fans as taking in Cassel's hit as just another element of the public spectacle of the game. This is the understanding that makes Winston's allusion to the Coliseum so fitting: a mass of people—a spectating crowd—treating another's pain as a piece of entrainment over which to gloat and commune together in shared joy.

you. It is the act of communication (and the attitudes thereby expressed) more than the belief in question, that makes her remark an apt target for blame.³ The blame, in other words, is exhausted by the communicative act.

Can we make a similar point about Winston's blame? Should we locate the blame's object solely in the fans' action (cheering)? I don't think so. I think Winston would be indignant at anybody who responded with happiness but failed to express it through cheering. I think Winston meant to blame those who were happy at the injury even if they refrained, *out of respect*, to communicate or express their emotion. I think he would just think of those people as more civilized assholes. And I think he would be right to do so. This is to say that wholly private attitudes—attitudes that are neither communicated nor expressed in behavior—seem to be worthy of blame. Winston's blame, even that directed at those who were happy but did not cheer, doesn't seem misdirected or confused or unfitting in any way.

The same point, I think, can be made by looking at excuses and justifications. Imagine a mother who had to make huge sacrifices to send her son to college. She had to sell her only car, use most of her savings, work overtime for years, etc. And she did this because she cares about him—that is, she did it *for his sake*. But suppose he mistakenly believes that she didn't pay much regard to what he really wanted. That she sent him to college not for his sake, but simply out of a vision of his future that, though he happened to share, this was entirely coincidental. And as a result he is not grateful for what she did for him. It's not as if he has said or done something to

³ A cleaner example of this would be a fan who wasn't really happy but nevertheless cheered. By hypothesis, such a person lacks the attitude his action 'expresses,' but he seems blameworthy for it all the same.

communicate his lack of feeling. He has always gone through the motions: he thanks her, says the right things, etc. And yet, he is not really grateful, and she knows it and he knows that she knows it. But now, at his graduation ceremony, it hits him that he was mistaken all along—he realizes that he was too self-absorbed to see the true motives behind her actions.⁴

In this sort of situation, I believe, it is natural for her son to offer an excuse or a justification (e.g., ‘I mistook your determination for a lack of interest in my view of things’). It is natural, that is, to offer a reinterpretation of the *significance* of the offense (cf. McKenna, 2012: 74-78). His excuse communicates that, contrary to what it may seem, it wasn’t out of ill-will or mere indifference that he did not feel gratitude, but out of a mistaken belief about her mother’s motives. But crucially, such a plea is not about his behavior, for by hypothesis, he didn’t behave objectionably in any way. It is about his feelings.

If there were no excuse, moreover, it would be an intelligible response to offer an apology. He might say, ‘I’m sorry. I shouldn’t have felt the way I did. I should have been grateful for all you have done.’ And of course, it would be an intelligible response from her part to grant him forgiveness, and eventually to forgive him. But the object of the apology—what he would be apologizing for—would not be an action or any other piece of behavior. He would apologize for not being grateful. And the same is true of her forgiveness.⁵

⁴ I owe this example to Steve Darwall.

⁵ One may argue that the object of the apology is not really the failure to *feel* gratitude, or of *being* grateful (if those two can come apart), but rather the mistaken judgment that she didn’t act for his own sake. Now, I don’t wish to deny that such a judgment may be part of his apology. It probably is. However, my sense is that the emotion itself is crucial—it is his failure to be grateful that would keep him up at night. We can see this more clearly, I believe, by imagining a case in which he makes the right

One could still object that it is pointless to argue that it is the attitudes themselves, regardless of their expression, that we care about in these examples. For it might be noted that we always experience others' emotions through some act or gesture that expresses them. Their individual impact, the objector concludes, can't therefore be isolated in the way I'm claiming.

The objection misses the point of my claim. The point is that when certain attitudes are expressed, either in deed or in words, we respond not merely to their *expression* but also, and often primarily, to their *existence*.⁶ It is not only our partner's petty remark that bothers us. It is *his being* contemptuous that does. Sure, the fans' cheering is disrespectful, but what really gets to Winston is the fact that Cassel's injury made them happy. To be clear, I'm not saying that actions are important only to the extent that they bring forth certain attitudes. Actions are not simply the medium through which the message is expressed. All I'm arguing here is that the mere holding of certain attitudes—and not only their behavioral expression—is enough to engage our accountability practices.

judgments but still fails to be grateful. He recognizes the sacrifices she made for him, acknowledges that she made them for his sake, but is left cold by the recognition. This case shows that the judgment is not enough. So I don't think we can reduce obligations of feeling (if there are any) to obligations to make certain judgments (or to have certain beliefs). See Basu (2019) and Marušić & White (2018) for the latter kind of obligations.

⁶ The point is obvious in first-person cases. We often feel guilt for having an emotion prior to even having the possibility of expressing it. See Smith (2011) for an insightful defense of the rationality of this kind of guilt. Smith theorizes the existence of "attitudinal obligations" and argues that we can make sense of them under a contractualist framework.

Suppose, then, that I have convinced you that our blaming attitudes often target others' emotions whether or not they are expressed. And suppose you are convinced that such responses are not confused or misguided but are sometimes fitting. What follows? Specifically, is this enough to conclude that our ordinary practices presuppose obligations of feeling—moral obligations to have and refrain from having certain attitudes? Or is the fact that emotions can't be demanded a sign that our practices carry no such presupposition?

Following Mill's first point, Stephen Darwall has given a conceptual analysis of the notion of a moral obligation. According to Darwall, "Necessarily, an act is wrong (violates a moral obligation) if, and only if, it is an act of a kind that it would be blameworthy to perform, where the agent to do so without excuse" (Darwall, 2017: 5).⁷ In other words, it is a conceptual truth that obligation is entailed by, and in turn entails, accountability. In a similar vein, Wallace has argued that the reactive attitudes, in particular resentment, indignation, and guilt involve the belief that their target has violated a moral obligation (see Wallace, 1994: esp. Ch. 2 and p. 77-8). According to Wallace, the belief figures both in the explanation of the occurrence of the attitude and in its representational content. Thus, both Darwall and Wallace draw a tight connection between what we are accountable for and what we are obligated to do.

⁷ Allan Gibbard gives a similar formulation, "what a person does is *morally wrong* if and only if it is rational for him to feel guilty for having done it, and for others to be angry at him for having done it" (Gibbard, 1990: 42, emphasis in the original).

And so, if we rightly hold each other accountable for our emotions, as I have argued, then our obligations are not—and could not—be restricted to outward behavior. Rather, they would have to encompass many aspects of our mental life.

There are, as I see it, two plausible strategies to get out of this conclusion. The first grants the point that we are often blameworthy for our emotions but denies that this shows that we hold each other accountable for them. The view denies, that is, that blaming is a form of holding accountable. T. M. Scanlon (2008), for example, has given a theory in which blame “is not, even incipiently, a form of communication” (fn. 54, 233-4). It is rather a modification of attitudes in light of an impairment to a given relationship. According to Scanlon:

[T]o claim that a person is *blameworthy* for an action is to claim that the action shows something about the agent’s attitudes toward others that impairs the relations that others can have with him or her. To *blame* a person is to judge him or her blameworthy and to take your relationship with him or her to be modified in a way that this judgment of impaired relations holds to be appropriate (128-9, emphasis in the original).

Scanlon himself acknowledges that attitudes may be in themselves blameworthy (157-8), but denies that this has any necessary connection with holding accountable, and hence with obligation. Blame marks an impairment in one’s relationship with the blamed, not the violation of an obligation.⁸

⁸ Scanlon emphasizes that one might impair a relationship without violating a moral obligation. He thinks this is especially true when it comes to unexpressed attitudes: “being a friend involves being disposed to certain feelings... A friend is not *obligated* to have such hopes and feelings, but a person who fails to have them, if a friend at all, is a deficient one” (132, emphasis in the original).

Even if Scanlon is right about the nature of blame, we still need a way to distinguish between impairments that are experienced as disappointing, perhaps even painful, on the one hand, and those that are experienced also as wrongful, on the other.⁹ My friend's malicious joy at my divorce impairs our relationship. But so does his newfound devotion for conspiracy theories. Though they might be equally threatening to the existence of the relation, only the former is experienced as a wrong—as a lack of regard. And so there would be something fundamentally different in my responses to the two cases, even if Scanlon wants to call both a form of blame. My blaming response to the former, but not to the latter, would seek to hold him accountable. In blaming him, for example, I might expect an apology or at least an acknowledgment of his lack of regard. None of this holds with respect to his conspiratorial tendencies. The point is thus the following: even if not every form of blame is a form of holding accountable, the sort of blame that I claim is often directed at private attitudes is.

A different strategy, therefore, grants that we hold each other accountable for our emotions but denies that this implies that our obligations extend to them. The idea here is to resist the conceptual connection between accountability and obligation. The scope of responsibility as accountability, in other words, is not limited to violations of obligation. On this view, one may be held accountable for something which one is under no obligation to do or feel.

Michael McKenna (2012), for example, agrees that in the prototypical case one is held accountable for violating an obligation. But he argues that there is a loosened

⁹ See Shoemaker (2011) and Smith (2013) for different elaborations of this thought.

sense of accountability, one without the “presumption of entitlement of the one blaming to blame” in which

the burdens of blaming are lessened, and so with it the harm that blaming in this fashion is likely to cause. Hence, the blaming, as in comparison with when an agent has violated an obligation... will pose less of a threat to the blameworthy person...Indeed, in many cases, it is likely that there will not even be an expectation that in blaming, the agent needs to change anything about how she *conducts* herself in her modes of dealing with others (202, emphasis in the original).

In this passage, McKenna is referring specifically to accountability for the non-voluntary, and so presumably to the type of accountability that is involved in the case of emotions. But I see no reason to think that this is the form that accountability has to take in such a case. On the contrary, if we look at Winston’s reaction to those “who even smiled,” it is not only clear that he felt entitled to blame them, but also that his blame carried the expectation that the fans should conduct themselves differently (where ‘conduct’ is understood here as encompassing the fans’ attitudes). I don’t think, moreover, that blame for attitudes is necessarily ‘low stakes’ or milder in comparison to blame for voluntary actions. This is especially evident in the context of personal relationships (though Winston’s case shows that it may be true more generally). A wife’s blame towards her husband’s unjustified jealousy can be as forceful (and as damaging) as her blame for any action he might do.

McKenna may be right that there is a distinction between accountability for obligations and accountability for other non-deontic stuff. But he is mistaken in

thinking that nonvoluntary attitudes (such as emotional responses) must fall on one side of this distinction.

My view is rather that accountability in its paradigmatic form, full-blooded accountability, is often directed at unexpressed attitudes. This suggests that we take ourselves to have obligations of feeling. That is, I take the fact that one can be held accountable for one's feelings as pretty good evidence that our obligations are not limited to actions. And this remains true even if there is no conceptual entailment from accountability to obligation.

4

But perhaps one cannot be held accountable for one's feelings. At least not without incurring a serious normative mistake. To posit obligations of feeling, the objection goes, is to *mislocate* morality—to fail to identify its proper place. The idea here is not to challenge the link between blame and obligation, but rather to claim that blame for unexpressed emotions is never justified. The mind is, in the words of George Sher, a “lawless wild west.” Sher (2019) writes,

[T]he purely mental is best regarded as a morality-free zone. Within that realm, no thoughts or attitudes are either forbidden or required. Unlike actions in the world, which morality is properly said to constrain, each person's subjectivity is a limitless wild west in which absolutely everything is permitted (484).

The proposal is that morality has nothing to say about our attitudes. Morality does not demand that we feel this or that. It demands, exclusively, to do some things and refrain from doing others—it is only what we do with our bodies that ought to conform to its rules. On Sher's view, then, the *mind* marks the proper limit of morality's demands.

But what is it about the mind in virtue of which it should serve as the relevant boundary? Suppose I see Cassel's injury and think, *Oh, Sweet!* According to Sher, we are all permitted to do this without morality stepping in. But suppose I'm alone in my room watching the game and instead of cheering inwardly I let out a little cry, 'Oh, Sweet!' It seems arbitrary to claim that morality cares about the latter but not about the former.

A better approach is thus to treat the cases in the same way. Sher could claim that neither your thoughts nor your spoken words fall within the moral domain. This would be to emphasize the *self*—not the mind—as the relevant limit. On this view, it is only when our attitudes and actions impinge on others that morality comes in. The idea here is that morality is essentially *public*. If my actions affect nobody—say I cheer Cassel's injury alone in my room with no intention that anybody ever hears me—then they are outside the jurisdiction of morality.¹⁰ It makes no meaningful difference whether or not I move my muscles—whether I yell at the television or say the words in my mind or write them down in my super-secret diary. If my actions are *private*, if

¹⁰ Of course, I don't mean to suggest that if my actions *just happen* to affect nobody then all is jolly good. I meant to refer to actions that are essentially private and not meant to involve others in any way. If one *mistakenly* thinks that one is acting privately when one isn't then one would *mistakenly* think one's action is outside morality when it isn't.

they involve myself and no one else, as far as morality is concerned, I am as good as new.

I take it that this reformulation (i.e., to posit the self rather than the mind as the relevant limit) remains true to the spirit, if perhaps not the letter, of Sher's view. For one, Sher's objection appears to remain intact. If morality's "central rationale lies in its ability to regulate our public interactions" (487), then our mental states, it would seem, cannot be immoral. Morality may demand actions that themselves require a given mental state to be carried out. It might be, for example, that a feeling of compassion is psychologically necessary to heed the demands of beneficent action. But this is quite different from demanding the mental state itself. That morality does not do.

The objection fails by its own lights. In other words, even if Sher is right about the public character of morality, he is mistaken about what should count as public. Mental states, I believe, can fail to be private in the relevant sense. Suppose I use my computer to photoshop a picture of my friend and create a pornographic image for my own personal use (I make sure that no one ever sees it or finds out about it). Is this action private? Yes and no. But 'no' in the sense that matters here. After all, it is *her* that I see when I look at the picture. It is an image of *her* body that I am manipulating for my enjoyment. Now suppose that instead of using the computer to manipulate the picture, I simply use my mind. Does it now seem any less absurd to suggest that my mental picture of *her* does not concern her, that it is mine in a way that stands outside morality?

The key point is that these attitudes involve other people, and so, morally, they are not “gloriously and completely our own” (495). If I’m happy about Cassel’s injury, then my happiness is *about* Cassel. It has *his* pain as its object. My emotion refers to another person and represents something about him in relation to myself. In this case, roughly, my pleasure represents *his* injury as something good for *me*.

But, someone might complain, as long as Cassel (or my friend) does not find out, he won’t be affected or harmed or have his interests interfered with in any way. Maybe. But that does not show that my pleasure about his injury is none of his business. The emotion already implicates him. It is private not in the deep sense that it is something that does not concern him but only in the thin sense that it is something he lacks access to.

If morality is about our relation with other people, as Sher claims, and I agree, then our emotions about other people—the way we respond, say, to their pain and suffering—must fall within its purvey.¹¹

¹¹ Sher presents a dilemma for those who defend the view that all kinds of mental activity, such as fantasies and imaginings, are subject to moral assessment. He argues as follows: take one of the mental states under dispute (e.g., a malicious fantasy) and ask, *Who is morality protecting? Who is being wronged by the attitude?* There are two options. One option is to claim that it is the representation in the mind that is wronged by my fantasy. But this is absurd. As Sher puts it, “The ‘people’ who populate our mental landscapes are only shadow people, and you can’t have a moral obligation—even a shadow obligation—to a shadow” (494). The other option, clearly more promising, is to say that the obligation is directed to the person that the fantasy is about—the person in the real world. The problem, Sher argues, is that many of these fantasies are not about an actual person. The second option therefore fails to explain what’s wrong with those mental states. This is a powerful argument, but it is not an argument against my position. For my argument concerns *exclusively*, and importantly, those mental states that do have as their object a real person *and* are a response to something real (Cassel’s pain, the son’s mother’s sacrifice, my friend’s body).

If blame is justified, and the presence of blame suggests the presence of an obligation, why do demands seem so absurd? What needs explaining, in other words, is Mill's second point. If I really have a duty of feeling, how come no one can demand that I fulfill it? One might think, for instance, that the fans are wronging Cassel for enjoying the vicious hit that knocked him cold. And yet, there would be something *disconcerting* were he to demand that they stop feeling happy. How to explain this?

Here's one obvious answer: demanding an emotion is infelicitous because it is not under the voluntary control of the person to whom the demand is issued to comply. And demands, *qua* speech acts, aim at compliance. As Macnamara (2013b) puts it, "A demand is successful as the kind of thing it is only if its target does as directed because she was directed" (897). But this is simply not possible when it comes to our emotions, for "[w]e can no more feel an emotion on command than we can digest our food on command" (902). Demanding an emotion, then, is like demanding someone to digest their food. It makes no sense.

This is no doubt a powerful argument. But I think it is mistaken, for it fails to identify the true source of the confusion. Now, I don't wish to deny the point that a successful demand is one in which the target does as demanded *because* it was demanded. Nor do I wish to deny Macnamara's claim that we can't do that in the case of emotions.¹² My point is rather that voluntary control is a red herring. I believe that the infelicity at issue would remain even if we had the relevant control over our

¹² Philosophers sympathetic to the idea that there are obligations of feeling have often tried to contend that we do in fact have the necessary control over our emotions. See, e.g., Liao (2006). I wish to remain neutral on this point for the sake of argument.

emotions. To put it simply, the problem is that demanding an emotion is self-defeating. And this has nothing to do with the fact that we lack control over them.

On my view, the infelicity in demanding an emotion¹³ can be modeled on the infelicity involved in demanding a compliment. It is part of the nature of compliments that they come unbidden, at least insofar as they aim to express admiration or esteem. The success of its internal goal (and oftentimes the speaker's goal) depends on their target not asking for them, let alone demanding them. To demand a compliment, therefore, is to squash a condition for its very existence. It is, in other words, self-defeating.

Notice, however, that nothing said here involves the notion of control or voluntariness. Though it is perfectly within our control to give a compliment, and so in principle one can pay a compliment because one has been commanded to, the target would not really obtain what compliments are meant to offer, namely admiration or esteem. Instead, one's words would merely be an acknowledgment of the target's authority (or sheer power) to make such a demand. For the compliment to achieve its purpose, it has to arise (or at least has to have the appearance of arising) from recognition of an admirable quality in the target. Hence, to give a demand would be to give the other a reason that, if followed, would prevent one from getting what is being demanded. This is the sense in which demanding a compliment is self-defeating.

Now come back to the case of the ungrateful son. It might be true that he can't make himself feel gratitude on command. But what would change if he could? Would a demand now be felicitous? No. What the mother wants is her son to be grateful

¹³ I come back to Cassel's case below, where the relevant demand is not about having an emotion but about not having one.

because of what she's done for him—not because she demands that he be so. And therefore to demand that he be grateful is to give him a reason that, if followed, would prevent her from getting what is being demanded. Her demand would be self-defeating not in the sense that his emotion would not count as *gratitude*, for we are supposing that he can make himself grateful, actually grateful, on command. The demand is self-defeating because, if successful, she won't get what she wanted. If I'm right, however, this is not only a matter of wants. If I'm right, the son has the obligation to be grateful. And so the claim is that a successful demand would prevent him from giving her what he is obligated to give her.

But what is it that he is obligated to give her? What is the content of the obligation? The suggestion, already emphasized, is that the obligation is not to be grateful period. But rather, to put it in the terms of Barry Maguire (2018), to be grateful for the considerations that make gratitude fitting. The obligation, to be clear, is *not* to be grateful because gratitude would be *fitting* (given the sacrifices). After all, there are lots of cases where gratitude is fitting but we are not obligated to feel it. The obligation is to be grateful because of *the sacrifices* themselves.¹⁴

Can the mother demand *this*? Consider the following: 'Be grateful for what I've done!' Is it possible to comply with such a demand? To do so, it seems, one would have to feel gratitude for the fit-making considerations *while* doing so because it was

¹⁴ Maguire argues that there are no reasons of the 'right kind' for attitudes (see D'Arms & Jacobson, 2000 and Rabinowicz & Rønnow-Rasmussen, 2004). There are only considerations that make attitudes fitting. Now, my view is that the son has a reason to feel gratitude, but it is a *moral* reason. He has a moral reason to have an attitude for the considerations that make the attitude fitting—not because it is fitting, and not because he is obligated to have it, but rather for the first-order facts (namely, his mother's sacrifices) that make the attitude fitting. Maguire is explicit that his argument is restricted to reasons of the right kind (780), and he explicitly mentions moral reasons as prominent examples of wrong kind of reasons (786-87). So I don't take myself to be disagreeing with Maguire in any way.

demanded. There is thus an internal tension within the demand itself. Now, the mother can say, 'You should be grateful for what I've done.' This remark does not make a demand, but rather serves an epistemic function: it points to (or calls attention to) the fit-making considerations he appears to be ignoring.¹⁵ But perhaps one can turn this epistemic 'pointing to' into a demand. In other words, it might be that the following makes sense: 'Look at what I've done for you!' Here, the mother does make a demand. But it isn't a demand for gratitude. What she demands is for him to consider the first order facts—to redirect his attention. His feelings, hopefully, will follow. But if they do, they will come from his recognition of the facts (and not from heeding a demand). As such, complying with the demand is consistent with gratitude arising from the considerations that make it fitting. Still, it remains true that, much like compliments, the emotion itself has to come unbidden. Gratitude has to come 'from him.'¹⁶

Compare the case with that of a promise. In the latter case, the promisee can demand that the promisor discharge her duty. Such a demand, as Darwall has noted, gives the other a reason distinct from the reasons issuing from a duty to keep one's promise, namely, a second-personal reason: a reason that comes, and is inseparable

¹⁵ Of course, the statement might have all sorts of pragmatic effects. It might be heard as a complaint (consider: 'You should have let me know you were running late.'). My reading in the text comes out more clearly when uttered by a third party: the remark 'You should be grateful for what your mother has done' is most naturally heard as directing the son to the relevant facts (and not, say, as making a demand or issuing a complaint).

¹⁶ What if he could take a pill that not only makes him feel gratitude, but makes him grateful for what his mother has done for him? If he takes the pill, that is, he will be able to recognize the relevant reasons and respond to them appropriately. Would she get what she is owed? Would she be grateful for her son's gratitude? The answer would depend on how we conceive of the drug's effects. If we see the drug as simply pushing neural buttons to create an 'artificial' emotion, the answer would be in the negative. If, on the other hand, we see it as enabling his true self to come out—as making it possible for him to express himself—then I see no reason to think of his gratitude as any less meaningful. In the latter case, it might even make sense to demand that he take his meds.

from, the demand itself.¹⁷ In fulfilling the obligation, it makes sense to act on that reason: it makes sense to act because the promisee demands it (Darwall, 2006).

In the case of gratitude, however, there is no comparable second-personal reason because the comparable demand would be self-defeating. It would give a reason that if responded to would make the obligation go unmet. The demand thus gets in the way of feeling what one is obligated to feel.

One might think, contrary to what I've argued, that the infelicity in question is not due to the general logic of demands, but rather to something unique to cases in which the person issuing the demand also figures in the content of the emotion being demanded. In the case of gratitude, for example, it could be that, in those cases, a demand changes the balance of (moral) reasons the target has to be grateful or at least the considerations that make gratitude fitting. Such a demand would not leave everything else as it was. Instead, what was fitting or appropriate or obligatory *before* the demand may no longer be once the demand has been made. The mother's demand for gratitude might cast her previous actions in a new light. It might suggest, perhaps, that what she wanted was his gratitude and not his wellbeing.¹⁸ The result is that gratitude may no longer be fitting or morally obligatory (or neither).

Now, I don't wish to deny that this might be true in some cases. I think it is. But this explanation fails to diagnose the infelicity we are trying to elucidate. To see this, consider third-party cases. Suppose it is now a friend who issues the demand: 'Be

¹⁷ Darwall defines second-personal reasons as follows: "What makes a reason second-personal is that it is grounded in (*de jure*) authority relations that an addresser takes to hold between him and his addressee" (Darwall, 2006: 4).

¹⁸ There's an interesting question here about the mechanism at play. Is it that the demand plays an epistemic function, *revealing* that gratitude is after all not required (i.e., should we take the demand as *evidence* that gratitude was never required)? Or is it that the demand itself *modifies* the normative landscape, altering the balance of reasons in such a way as to defeat the obligation to have the feeling?

grateful to your mother!’ Here, the demand can’t possibly modify or call into question the normative significance of the *mother’s* actions. And yet, complying with the demand would not give her what she wants (and what the friend wants her to have), for feeling gratitude because *he* demands it is no better than doing so because *she* demands it. Both are self-defeating in the same way.

The same dynamic ensues in actions aimed at expressing an attitude. To demand that you call me on my birthday is self-defeating. It undermines the very point of birthday calls. Can others make that demand of you? Sure. In the typical case your spouse will say: ‘Call Mario, it’s his birthday!’ This is all good...as long as I don’t know about the demand. The more I suspect that that was your reason, the less I would appreciate the call. At the limit, your call is worthless. It shows, simply, that you are an obedient spouse. This gets really absurd really quickly. My own grandmother religiously reminds me every year of my parents’ birthdays a day in advance. My parents, of course, know this, since she reminds everyone in the family of everyone else’s birthdays. So my calls—and those of everyone else in the family—lack the expressive meaning they would otherwise have. And this is true even if I would remember their birthdays anyways. In my family, a birthday call is forever ambiguous: it can express appreciation, sure, but more often than not it simply reflects compliance with my grandma’s (subtle) demand. Her reminders are nice, but self-defeating nonetheless.

This shows—conclusively, I believe—that the alternative explanation can’t be the whole story. If it were, third-party cases would not present any problem. But they

do. And the problem is the same: to demand gratitude (and whatever attitude is expressed in birthday calls) is self-defeating regardless of who makes the demand.

Let us now turn to a different question. Is it always self-defeating to make demands on our attitudes? Or are there contexts in which no such problem arises? Suppose Cassel were to say, 'Stop being a dick!' or 'Stop rejoicing in my pain!' What should we say here? On the one hand, there *is* something odd in such a demand. One usually wants and expects others' reactions to one's suffering to come unbidden. Cassel presumably wants the fans to react appropriately not because he demands it but simply because he is in pain. To the extent that this is Cassel's motive, the demand is in fact self-defeating. On the other hand, Cassel may simply want the fans to show some respect. And respect is something that can result from heeding a demand. To see this, consider someone who tells her boyfriend, 'Stop lusting after other women!' Here, the demand successfully gives a second-personal reason in Darwall's sense. She wants him to stop lusting after other women *out of respect for her*.¹⁹ The boyfriend acknowledges her authority to make such a demand by doing as demanded because it was demanded. Similarly, the fans may acknowledge Cassel's authority to make such a demand by refraining from indulging in their joy.

This shows, I believe, that there are felicitous demands we can make on our attitudes. And as we have seen, there are actions that are self-defeating in the way demanding gratitude is. The relevant contrast, then, is not between actions and attitudes. Nor is it between what is under our control and what isn't. It is between

¹⁹ And not, say, out of his finding other women objectively unattractive. But now suppose she comes back and says, 'and start lusting after me!' The demand is funny because absurd, for presumably she wants his boyfriend to find her attractive on his own, and not because she demands it. As Darwall says in the case of love, "Love is nothing that could be given as a result of accepting a valid claim; love cannot arise from respect" (Darwall, 2016: 99).

what can be successfully given as a result of heeding a demand and what cannot be given in that way.

Finally, it is worth remarking that it is a mistake to infer from the fact that an action or an attitude falls on the latter class that one is under no obligation to act or hold that attitude. That it would be self-defeating to demand a birthday call, in other words, does not mean that there is no obligation to call.²⁰ What follows is only that it would make no sense to demand it.

6

I seem to have arrived at an awkward position. If demanding an emotion is self-defeating how come morality does it? The argument I just gave, that is, appears to apply to demanding an emotion as such, irrespective of who or what is making the demand. You don't want your son to be grateful because you demand it. But do you want him to be so because morality demands it? The latter does not seem any better than the former. So if I'm right and demanding that he be grateful is self-defeating, morality's demand would be self-defeating too. In each case, the obligation would not be discharged if the demand were successful. Morality appears to undermine itself by making a demand that can't succeed *by its own light*—can't be followed, that is, without failing to deliver the very thing being demanded.

²⁰ Here again the comparison with compliments is useful. The fact that demanding a compliment is self-defeating in no way bears on the question of whether paying a compliment would be deserved or fitting.

Doesn't this show that morality can't make such demands and therefore that no demands of this kind exist? Doesn't this reveal an inherent tension in the very concept of an obligation of feeling, a tension that should lead us to reject the concept?

The first thing to recognize is that the language of *morality's demands* is figurative, and confusion arises by taking it too literally. To see this, recall that it is in the nature of demands to aim not only at compliance, but also at compliance for the specific reason that a demand has been issued. A successful demand is one in which one does as demanded because it was demanded. But no one really thinks this in the case of the so-called demands of morality. Suppose morality demands that we give 10% of our income to the very poor. And suppose you do so because you care about their plight. It is uncontroversial, I hope, that you have successfully fulfilled the demand. In terms of what morality demands, your action is in perfect compliance. And yet, you didn't do it *because morality demanded it*. Notice that this is so regardless of what we say about the moral *worth* of your action. Kantians regard the same action done from duty—i.e., done because morality demands it—as having higher worth. Others strongly disagree. But everyone, Kantians (and Kant) included, agree that as far as meeting morality's demands, you have done your due. So it is a mistake to think of morality's demands in the way we think about the speech act of making a demand. It's more useful, I believe, to talk about obligations or requirements, which need not figure in the content of the agent's motive to be successfully met.

Still, one might insist that a tension remains. The obligation to give to the poor is met whether or not the motive of my donation references the fact that I am obligated to give. Moreover, there is nothing objectionable or off-putting in answering

the question, 'Why did you give?' with 'Because it is my duty.' And yet, my claim is that this would be the wrong answer in the case of emotions. As I argued in Section 5, obligations of feeling, to be successfully discharged, often require that they be a response to the considerations that make the feeling fitting. And this seems to require the motive of duty to be absent. The objection then is that this is implausible. It is implausible to think that the motive to meet requirement *X* can prevent *X* from being met. The problem may not be one of self-defeat. But there would be some sort of instability in the requirement if morality were to behave like this.

I agree that there is some instability here. But it does not lie *in* morality or in what it requires, but rather in certain ways of *thinking* about them. Bernard Williams (1981) gives an example of this sort of "reflexive deformation." An action, he says, displays a person's integrity when it arises from "those motives which are more deeply his" (49). But integrity *as a motive* is suspect. To truly act with integrity, it seems, one can't have integrity itself as a motivating thought. That would be a "misdirection of the ethical attention," as he puts it elsewhere in Murdochian terms (Williams, 1985: 11). The point I want to make is simply this: the fact that integrity itself cannot be turned into a motive—the fact that being motivated *by* integrity makes it harder to act *with* integrity—does not mean that there is a problem with the notion of integrity. Or that we should not think that displaying integrity is a good thing. What it means is rather that "the importance of an ethical concept need not lie in its being itself an element of first-personal deliberation" (Williams, 1985: 11).

And so, similarly, there are obligations which are better fulfilled by not thinking about them in those terms. Some obligations may function like integrity.

They may be such that one cannot fulfill them (directly) by taking on the motive to do so. Consider the following famous example by Michael Stocker (1977):

[S]uppose you are in a hospital, recovering from a long illness. You are very bored and restless and at loose ends when Smith comes in once again. You are now convinced more than ever that he is a fine fellow and a real friend—taking so much time to cheer you up, traveling all the way across town, and so on. You are so effusive with your praise and thanks that he protests that he always tries to do what he thinks is his duty, what he thinks will be best. You at first think he is engaging in a polite form of self-deprecation, relieving the moral burden. But the more you two speak, the more clear it becomes that he was telling the literal truth: that it is not essentially because of you that he came to see you, not because you are friends, but because he thought it his duty (462).

Stocker was making a different point, but the example shows, I think, that taking on the motive of duty might interfere with the successful fulfillment of one's duties. For it does not seem like a stretch to suppose that Smith has fallen short of meeting the obligations that, *qua* friend, he held towards you.

One way in which this interference might happen is rather simple. The actions that one might undertake out of a motive of duty might differ from those undertaken out of, say, friendship or benevolence. Hence, one might fail to act rightly because one might fail to perform the required actions. But this does not seem to be the case with Smith. Rather, the problem here appears to be with the motive itself, for Smith's actions seem unimpeachable. And so, insofar as Smith fails to meet a requirement, it has to be that some requirements (perhaps friendship is paradigmatic in this respect) have the agent's motivation as part of their content (cf. Liao, 2006: 26). But if there are such requirements, it necessarily follows that they are subject to Williams-style

reflexive deformations. If Smith's duty is not only to visit you, but also to do so with the appropriate motivation—to be moved by the right reasons—then it's no wonder that his preoccupation with duty conceived as such may be an impediment in fulfilling it.

But notice that if obligations that have both a behavioral and an affective/motivational component are prone to reflexive deformation, then those which consist solely on the latter would be so as well. For it is at this level where the effects of reflexive thought occur. In the case of emotion, therefore, we should similarly expect that an explicit preoccupation with duty might deform or at least inhibit the motives and feelings the obligation specifies, thereby interfering with its fulfillment. But there is nothing questionable about this (nothing to make us question the existence of such obligations). Some goals are better achieved by not aiming directly at them. Pleasure is said to behave like this. So is happiness. Williams has shown that some character virtues exhibit this feature. My claim is that we can add a subset of moral obligations to the list.

Some might object as follows: 'We engage in moral inquiry not to discover some hidden truths but to change our conduct. We want to know what obligations there are so that we can act differently. Learning that one has a moral obligation has to have practical consequences. If I am told that I am obligated to *X*, then there is

something I can do with that knowledge, namely, *Xing*. To realize that I am obligated to give to the poor is to realize that I can do something differently, namely, giving to the poor. But what can I do if you tell me that I am obligated to be grateful to my parents? I can't feel gratitude on command. At the most, I can endeavor to become the sort of person who is grateful for what others have done for him. Or I can behave as if I'm grateful—I can do the sort of things or utter the sort of words that usually convey gratitude. But if you tell me that my obligation is to feel gratitude—to be now grateful—then what you say is worthless to me. There is literally nothing I can or could do with what you are telling me. Your thesis is useless. There are no obligations of feeling (or if there are nobody should care).'

The picture of obligation expressed in the objection is one many philosophers favor. In fact, the claim that obligations are action-guiding is usually at the front of many arguments in support of the '*ought*' implies '*can*' principle (see, e.g., Copp (2003); Margolis (1967); Sapontzis (1991); Vranas (2007)). I think this is a mistaken picture—or at least an incomplete one. What's missing is the obligee, the person to whom the obligation is directed at. An all-out, agent-centric view of obligation obscures the important point that the notion of obligation serves to mark the fact that someone has been or would be *wronged* by a given action or attitude. And this of course carries all sorts of implications, practical and otherwise.

Let me give an example. Suppose you are part of a trial jury in a case that will have life-defining consequences for many people. There is a piece of evidence that weighs highly on your estimation of the case—say you have prior knowledge of the criminal history of the defendant—but that the judge has rightly deemed inadmissible.

Assume that you have a moral obligation—not just a legal one—to ignore such evidence in your deliberation. But suppose that you are psychologically incapable of doing that. You can't figure out what you would believe were you not to know about the criminal history of the defendant.²¹ It is a cognitive exercise you simply can't perform successfully. So what can you do with the judge's reminder that you are obligated to ignore the evidence? Very little, I think. In fact, it may be that the more you try to fulfill your obligation, the more you think about the evidence and the harder it is to ignore it. The obligation is therefore useless (or worse) as an action-guiding reason. But does it follow that you are not in fact obligated to do so? I don't think so. I think it's clear that you are required to ignore the evidence in your deliberation even if you can't do that.

But what's the point of saying you have an obligation if you are not capable of fulfilling it (it is not under your voluntary control to do so)? The answer, I believe, lies in the defendant. Each member of the jury owes her a fair and impartial verdict, and therefore she would be wronged were you to come to a judgment taking her criminal history into account. By calling it an obligation we mark the fact that your biased judgment is not merely an unfortunate occurrence with which she has to deal, but a wrongful denial of what she is owed, and for which she can rightly hold you accountable. Calling it an obligation enables us to understand—and justify—all sorts of reactions (e.g., blame or a demand for compensation and even punishment) from the part of the defendant that would be puzzling otherwise.

²¹ Or, perhaps more realistically, whatever it is that you do 'figure out' is colored by that knowledge, even if you are unaware of it.

Why not say, someone might insist, that the obligation is to *try* to ignore the evidence, for trying is surely under one's voluntary control and so it's something one can do? You owe the defendant your best attempt at a fair verdict (something you can do) but not actual success (something you can't). I find this implausible. 'Hey, they tried' would not do as an excuse for a biased jury. The obligation is to ignore the defendant's criminal history in reaching a judgment—anything short of that is a violation of that obligation. This is evident, for example, by the fact that it would be natural to be consumed with anxiety afterwards about whether your decision was in fact free of the influence of the defendant's history.²² And, if sufficiently self-aware, for you to start feeling the pangs of guilt. This is true even if you know that you tried to ignore the evidence.²³

Finally, it seems even more absurd to claim that given your inability to ignore the evidence your obligation should be traced (to use the term in the moral responsibility literature) to a voluntary action or decision in the past which resulted in the present inability (i.e., you were obligated *then* to do *X* or to refrain from doing *Y* but you are not obligated *now* to ignore the evidence, for *now* you can't).²⁴ Or similarly, that your obligation *now* is to become the sort of person who would be able to give a fair and impartial verdict next time around. Both options wipe the obligee out of view in order to cling to a myopic picture of obligation as necessarily action-

²² This is a familiar feeling to anyone in charge of grading student papers. You know that all you can do is *try* not to be influenced by the student's academic record. But you also know that the obligation is not only to try but to succeed.

²³ It may be that a juror that genuinely tried is less blameworthy than another who simply didn't care. For an account of degrees of blameworthiness that takes into account effort, see Arpaly & Schroeder (2014) and Nelkin (2016).

²⁴ For appeals to tracing, see Fischer & Ravizza (1998); Levy (2005). For critiques, see McKenna (2008); Smith (2005; 2008); Vargas (2005).

guiding. What you do in the future, or what you did in the past, is of little concern to the defendant to whom your present obligation is directed at.

Let us now think about the case of emotions in light of this example. Take the fans enjoying the sight of Cassel being knocked down to the ground. And let us grant for the sake of argument they lack control over their enjoyment (they really can't help but to feel happy about his injury). To call attention to their obligation is practically useless. But, as was the case with the juror, we want to mark the fact that someone has been wronged. To say that the fans have an obligation is not to give them a reason they can act on, it is rather to say that Cassel is wronged²⁵ by their pleasure, and thus to see as justified (or fitting) accountability responses such as Winston's. Winston's angry blame and condemnatory words make sense only under the assumption that the fans' emotion is wrongful—that they have violated an obligation.

What about the often-argued position of locating the obligation somewhere in the past or in the future? Here is Peter Vranas: "if, when I utter "you ought to feel grateful to her", I mean that you ought to be experiencing gratitude *right now*, then it is plausible to say that I am not ascribing to you any obligation to do something. I may instead be blaming you for not having cultivated in the past a disposition to feel gratitude whenever appropriate" (Vranas, 2007: 175, emphasis in the original). I think you are ascribing an obligation to 'do' something, namely to feel gratitude *right now*. It is as implausible to think that Winston is blaming the fans for not having cultivated a disposition to not enjoy a player's pain, as it is to think that the defendant would

²⁵ This follows, of course, only if we assume, as we have been assuming, that the fans not only have an obligation to refrain from enjoying Cassel's injury, but also that the obligation is an obligation *to* Cassel. In other words, it is the *directed* nature of the obligation, and not simply the obligation *per se* that implies that Cassel has been wronged. Similar remarks apply for the juror case above. I thank Steve Darwall for pressing me on this point.

blame the juror for not having cultivated a disposition to ignore inadmissible evidence. And the case of gratitude is not different in that regard. (Though it may be that it makes more sense with gratitude to *also* blame the lack of past cultivation). Similar remarks apply, *mutatis mutandis*, for the claim that the fans' obligation is only to cultivate such a disposition in the future. Both proposals make the mistake of bypassing the obligee, thereby misplacing the object of the accountability response.

Much more can be said and has been said about this issue.²⁶ My goal has simply been to argue against the charge that it is useless and hence mistaken to ascribe obligations that can't play an action-guiding role. I hope I've shown that the implications of ascribing an obligation go beyond the practical consequences this has for the agent. It is therefore no argument against obligations of feeling that the agent is often in no position to do anything directly with them.

8

Let me conclude by pulling the threads together. I started out arguing that we hold each other accountable for what we feel. We blame others for their emotions. And we blame ourselves. We often apologize. And sometimes forgive. Our emotions,

²⁶ In particular, I have said nothing about another traditional argument for '*ought* implies *can*,' namely, that it is unfair or unfitting to blame an agent for failing to do what she can't do or for doing what she can't refrain from doing (see, e.g., Copp (2003); Stocker (1971)). (Even philosophers who reject '*ought* implies *can*' often hold '*blame* implies *can*,' see Korsgaard (2018)). My view is that the argument assumes a punitive theory of blame; a theory we have good reason to reject (though I won't defend the claim here). For views of blame that deny it an essentially punitive aspect, see Darwall (2006; 2018); Hieronymi (2004); Macnamara (2013a; 2013b); Scanlon (2008; 2013); Smith (2013); Talbert (2012).

that is, engage our accountability practices in all their various textures and colors. But there is something we don't do. We tend not to demand others to have (or to refrain from having) the feelings we nevertheless blame them if they don't (or do).

Why? I argued that this is not because morality does not care about our emotions, but rather because demanding them is pointless. If I have an obligation to feel gratitude, and you demand that I be grateful, then *complying* with your demand makes it impossible to comply with the obligation. A demand is therefore self-defeating.

But doesn't this imply that it is self-defeating for morality to demand emotions from us? No, for morality doesn't make demands in the sense that would make its 'demands' self-defeating. Morality does not aim at compliance in the way demands *qua* speech acts do. Nevertheless, it might be self-defeating for the agent to attempt to meet her obligations of feeling by understanding them as such—as *moral obligations*. To meet her moral obligations the agent may have to take morality out of view. Such reflective deformation, however, is no objection. It simply makes salient the fact that the importance of morality is not exhausted by its role in practical deliberation. The same answer was given to a different objection: how can there be obligations that fail to be action-guiding? In response, I redirected the attention away from the agent and towards the obligee. We need the notion of obligation to mark the fact that the latter has been *wronged* and therefore is not simply the victim of bad luck. I argued that this has important implications for how we understand the obligee's reactions and the ensuing relationship with the agent.

In a nutshell, my proposal is that morality requires not only certain ways of treating others, but also certain ways of feeling about them. Or better, that what we feel about others is part of treating them in the way morality requires. To modify a remark of Korsgaard's, the subject matter of morality is not what we should do, but how we should relate to one another. *How should we relate to one another?* If that's the question morality asks, I have offered a constraint in giving it an answer. Any plausible account would have to include our feelings and not merely our actions.

References

- Adams, R. M. (1985). Involuntary sins. *The Philosophical Review*, 94(1), 3-31.
- Arpaly, N., & Schroeder, T. (2014). *In Praise of Desire*. Oxford University Press.
- Basu, R. (2019). What we epistemically owe to each other. *Philosophical Studies*, 176(4), 915-931.
- Copp, D. (2003): 'Ought' implies 'can', blameworthiness, and the principle of alternate possibilities. In *Moral Responsibility and Alternative Possibilities: Essays on the importance of alternative possibilities*, ed. David Widerker and Michael McKenna. Ashgate: 265–299.
- D'arms, J., & Jacobson, D. (2000). The moralistic fallacy: on the 'appropriateness' of emotions. *Philosophical and Phenomenological Research*, 65-90.
- Darwall, S. (2006). *The Second-person Standpoint: Morality, respect, and accountability*. Harvard University Press.
- Darwall, S. (2016). Love's second-personal character: reciprocal holding, beholding and upholding. In *Love, Reason and Morality*, ed. Esther Engels Kroeker and Katrien Schaubroeck. Routledge: 119-135.
- Darwall, S. (2017). What are moral reasons? *The Amherst Lecture in Philosophy* 12, 1–24.
- Darwall, S. (2018). Contempt as an other-characterizing, "hierarchizing" attitude. In *The Moral Psychology of Contempt*, ed. Michelle Mason. Rowman and Littlefield: 193-216.
- Fischer, J.M. and Ravizza, M. (1998): *Responsibility and Control: A theory of moral responsibility*. Cambridge University Press.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings*. Harvard University Press.
- Hieronimi, P. (2004). The force and fairness of blame. *Philosophical Perspectives*, 18, 115-148.
- Korsgaard, C. (2018). *Fellow Creatures: Our obligations to the other animals*. Oxford University Press.
- Levy, N. (2005). The good, the bad and the blameworthy. *Journal of Ethics and Social Philosophy*, 1(2), 1-16.
- Liao, S. M. (2006). The idea of a duty to love. *The Journal of Value Inquiry*, 40(1), 1-22.
- Macnamara, C. (2013a). Taking demands out of blame. In *Blame: Its nature and norms*, ed. D. Justin Coates and Neal A. Tognazzini. Oxford University Press: 141-161.
- Macnamara, C. (2013b). "Screw you!" & "thank you." *Philosophical Studies*, 165(3), 893-914.
- Maguire, B. (2018). There are no reasons for affective attitudes. *Mind*, 127(507), 779-805.
- Margolis, J. (1967): One last time: 'ought' implies 'can.' *The Personalist* 48, 33–41.
- Marušić, B., & White, S. (2018). How can beliefs wrong?—a Strawsonian epistemology. *Philosophical Topics*, 46(1), 97-114.
- McKenna, M. (2008). Putting the lie on the control condition for moral responsibility. *Philosophical Studies*, 139(1), 29-37.
- McKenna, M. (2012). *Conversation and Responsibility*. Oxford University Press.

- Mill, J. S. (2001). *Utilitarianism*, ed. G. Sher. Hackett Publishing Company.
- Mitchell, H. (2012, October 8). Chiefs players rip their own fans for cheering Matt Cassel injury. *Los Angeles Times*. Retrieved from <https://www.latimes.com/sports/la-xpm-2012-oct-08-la-sp-sn-chiefs-fans-cheer-matt-cassel-injury-20121008-story.html>
- Nelkin, D. K. (2016). Difficulty and degrees of moral praiseworthiness and blameworthiness. *Noûs*, 50(2), 356-378.
- Rabinowicz, W., & Rønnow-Rasmussen, T. (2004). The strike of the demon: On fitting pro-attitudes and value. *Ethics*, 114(3), 391-423.
- Sapontzis, S. F. (1991). 'Ought' does imply 'can.' *The Southern Journal of Philosophy* 29, 383-393.
- Scanlon, T. M. (2008). *Moral Dimensions*. Harvard University Press.
- Scanlon, T. (2013). Interpreting blame. In *Blame: Its nature and norms*, ed. D. Justin Coates and Neal A. Tognazzini. Oxford University Press: 84-99.
- Sher, G. (2019). A wild west of the mind. *Australasian Journal of Philosophy*, 97(3), 483-496.
- Shoemaker, D. (2011). Attributability, answerability, and accountability: Toward a wider theory of moral responsibility. *Ethics*, 121(3), 602-632.
- Smith, A. M. (2005). Responsibility for attitudes: Activity and passivity in mental life. *Ethics*, 115(2), 236-271.
- Smith, A. M. (2007). On being responsible and holding responsible. *The Journal of Ethics*, 11(4), 465-484.
- Smith, A. M. (2008). Control, responsibility, and moral assessment. *Philosophical Studies*, 138(3), 367-392.
- Smith, A. M. (2011). Guilty thoughts. In *Morality and the Emotions*, ed. Carla Bagnoli. Oxford University Press: 235-56.
- Smith, A. M. (2013). Moral blame and moral protest. In *Blame: Its nature and norms*, ed. D. Justin Coates and Neal A. Tognazzini. Oxford University Press: 27-48.
- Stocker, M. (1971). 'Ought' and 'can.' *Australasian Journal of Philosophy*, 49(3), 303-316.
- Stocker, M. (1977). The schizophrenia of modern ethical theories. *The Journal of Philosophy*, 73(14), 453-466.
- Strachan, Y. (2012, October 9). Chiefs Eric Winston rips hometown fans in cheering QB injury. *Digital Journal*. Retrieved from <http://www.digitaljournal.com/article/334476>.
- Strawson, P. F. (1982). Freedom and resentment. In *Free Will*, ed. Gary Watson. Oxford University Press: 59-80.
- Talbert, M. (2012). Moral competence, moral blame, and protest. *The Journal of Ethics*, 16(1), 89-109.
- Vargas, M. (2005). The trouble with tracing. *Midwest Studies in Philosophy* 29, 269-291.
- Vranas, P. B. (2007). I ought, therefore I can. *Philosophical Studies*, 136(2), 167-216.
- Wallace, R. J. (1994). *Responsibility and the Moral Sentiments*. Harvard University Press.
- Wallace, R. J. (2012, June). Duties of love. *Aristotelian Society Supplementary Volume* (Vol. 86, No. 1, pp. 175-198). Blackwell Publishing Ltd.

Watson, G. (1996). Two faces of responsibility. *Philosophical Topics*, 24(2), 227-248.

Williams, B. (1981) Utilitarianism and moral self-indulgence. In *Moral Luck: Philosophical papers 1973-1980*. Cambridge University Press.

Williams, B. (1985). *Ethics and the Limits of Philosophy*. Routledge.