Canadian Journal of Philosophy, 2015
Vol. 45, No. 1, 58–84, http://dx.doi.org/10.1080/00455091.2014.997334

Routledge
Taylor & Francis Group

Mere moral failure

Julie Tannenbaum*

Department of Philosophy, Pomona College, 333 North College Way, Claremont, CA 91711, USA

(*Received 10 June 2014; accepted 18 November 2014*)

When, in spite of our good intentions, we fail to meet our obligations to others, it is important that we have the correct theoretical description of what has happened so that mutual understanding and the right sort of social repair can occur. Consider an agent who promises to help pick a friend up from the airport. She takes the freeway, forgetting that it is under construction. After a long wait, the friend takes an expensive taxi ride home. Most theorists and non-theorists react to such cases by either judging the agent's action as a violation of her obligation to help or as having satisfied the only obligation she really had, namely to try to help. However, as I show, there are serious difficulties that arise from categorizing this agent's action as satisfying or violating her obligation – difficulties that are avoided if we instead add "mere moral failures" to the basic categories for moral evaluation. An agent merely fails when she neither satisfies nor violates her obligation. She is responsible for what she has done, yet without thereby having done wrong. Moreover, there is a recognizable (though nameless) reactive attitude reflecting the agent's responsibility that falls between blame and mere regret, and is also importantly different from agent-regret. What I show is that we need not reach for blame as our only way of registering that an agent has morally failed another. While thus far overlooked, mere moral failure is by no means a rare occurrence, but rather a regular part of life among friends, family, investors and clients, police and citizens, doctors and patients, and many others.

Keywords: moral obligation; violation, wrongdoing; blame; failure; agent-regret

Introduction

Every day there are things we are supposed to do (or not do); sometimes we manage it, and other times we fall short. Ethical theorists use various terminologies to capture this intuitive idea of a requirement to act: we have an obligation to act, we should act, we ought to act, and so on.[1] When we do what we are supposed to do, we meet or satisfy our obligation, do what we ought, or do the right thing. However, if the window of opportunity to act has passed and what one was supposed to do has gone undone, then it is taken as obvious that not only has one not satisfied the obligation, but one has also violated the obligation or done what is wrong.

*Email: julie.tannenbaum@pomona.edu

© 2015 Canadian Journal of Philosophy

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

While the terminology might differ, what both theorists and non-theorists seem to embrace is a dichotomy of categorization that arises when we are faced with an obligation to act. Once the opportunity to act has been seized or gone undone, satisfaction and violation exhaust the possibilities for morally evaluating such actions.[2] However, as I will show, there are serious challenges facing those who accept this dichotomy – difficulties that can be avoided if we instead add "mere moral failures" to the basic categories for moral evaluation. An agent merely fails when she neither satisfies nor violates her obligation.[3]

Let us begin with some cases that illustrate the variety of actions that (so I will argue) fall into the category of mere moral failure.

> One neighbour promises to help out another by picking her up at the airport at 6 pm and she heads out the door in plenty of time to do so. She has no habitual way of getting to the airport; this time she decides on the freeway. As she enters a thicket of unusual traffic, she remembers yesterday's announcement about the start of freeway construction. By the time she gets off the freeway and takes side streets to the airport she is so late that her neighbour (who does not have a cell phone and so cannot be reached), after a long wait, finally takes an expensive taxi ride home.

> Strolling to the end of the pier, a person notices a stranger struggling in the ocean. No one else is around. She cannot swim but she sees a nearby rope. She throws one end of the rope at the stranger, but the currents carry it immediately out of his reach. She cries out, "Of course, the currents!" and taking the currents into account re-throws the rope. But just then, the stranger drowns.

> A friend is holding my child as he walks into a restaurant. He notices that the floor is wet and takes care to step where it is dry, so as to avoid slipping and dropping my child. However, he forgets to look up, does not notice a low hanging lamp, and so bangs my child's head into the lamp as he moves forward.[4]

The first two agents have an obligation to aid, however different the sources of obligation may be, while the third has an obligation not to harm.[5] As I will argue (Section 1), their lapse in recalling the relevant fact(s) does not undermine the claim that they have these obligations.

Those in the grip of the dichotomy might react in a variety of ways, which I discuss in the coming sections; one such response might be: these agents are morally responsible; they did not try hard enough or take enough care; they are negligent; they have acted wrongly and blaming them is justified.[6] Suppose I then point out how hard these agents were trying, precisely because they cared, that precautions can be foiled by one-off memory lapses that are not in any way blameworthy and so on. Those in the grip of the dichotomy might then flip to maintaining (or have already thought from the start) that there is no wrongdoing in these cases, and hence that neither responsibility nor blame is called for. These agents are unlucky and imperfect, but since they were not negligent, they are only justified in feeling sadness, mere regret, and disappointment at having made a costly mistake, but not guilt, and others are justified only in feeling frustration and wishing that things had gone otherwise, but not in feeling resentment. However, this second response does not do justice to the lingering thought that

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

these agents are nevertheless somehow on the moral hook; they are morally responsible and they need to make it up to the other person (if possible).

These cases, and others like them, give rise to two seemingly conflicting reactions. On the one hand, (a) these agents are on the moral hook – they are morally responsible and have morally failed the other person in some way – and an emotional reaction that reflects this is justified. On the other hand, (b) a specific set of emotional reactions is out of place: guilt (first personal reaction), resentment (second personal reaction) and indignation (third personal reaction). That is, blame, using it as shorthand for all three types of reactive attitudes, is not justified.

I plan to take seriously the phenomena as described by (a) and (b), i.e. accept that these agents are morally responsible, and that the feeling that they have failed another is justified, but blame is not justified. Relying on (a) and (b), I present a variety of challenges to those who maintain that the dichotomy (violation vs. satisfaction) is exhaustive (Section 2).[7] My goal is to show that however we variously carve up the moral landscape with the categories and concepts already in our arsenal (e.g. obligation, satisfaction, violation and blame), we will misdescribe and distort our understanding of the above-mentioned cases unless we make room for the category of mere failure and its distinctive accompanying reactive attitude.

While there is likely more than one way to accommodate the category of mere moral failure, I suggest that one way to construe the conceptual landscape is this (Section 3): some agents satisfy their moral obligation to act, while others fail to do so, and of those who fail, some merely morally fail and others violate their obligation (do wrong).[8] The category of mere failures fits nicely with (a) and (b) and shows that they can both be held without conflict. Those who fail their obligations are morally responsible, but one can be responsible without thereby being to blame. We need not reach for blame as our only way of registering that an agent has morally failed another. The category of mere moral failure also has interesting implications for determining the necessary and sufficient conditions for satisfying an obligation to act.

Upon discovering a conflict between phenomena like (a) and (b) and theory, some philosophers prefer to explain the phenomena away rather than adjust their theory. I am not arguing against that strategy, but instead working within a tradition, which began with Aristotle and continues in the work of more recent philosophers like Peter Strawson and Bernard Williams, that gives priority to the certain kinds of emotional phenomena.[9] Thus, this article presupposes (a) and (b). In Section 2.1, the above-mentioned cases are contrasted with those where blame is justified, while in Sections 1 and 2.2 the contrast is drawn with agents who are off the moral hook of responsibility. These contrasts may help some readers who do not yet share the reactions (a) and (b) move closer to those reactions.

The category of mere moral failure and its accompanying reactive attitude has been overlooked in the philosophical literature. And yet, mere moral failure is by

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

no means a rare occurrence; it is, rather, a regular part of life for finite rational beings such as ourselves. Given how often we are merely failed and fail others, in spite of good intentions, it is important that we have the correct theoretical description of what has happened, which requires expanding our moral theorizing beyond that of wrongdoing (violation). These unfortunate relations arise among friends, family, investors and clients, police and citizens, and doctors and patients, to name just a few examples. It can be tempting, especially for the injured or their loved ones, to jump to blame. And yet when the failure is like that of the three above-described cases, then neither blame, detachment, nor mere sorrow accurately emotionally reflects what has happened. If there is to be any hope of self and mutual understanding, acceptance, psychological healing, and the right sort of social repair, it is crucial that we understand what our obligations to others are and what it really takes to satisfy (or violate, merely fail, etc.) those obligations.[10] These issues come up not only for those who fail and are failed, but also for their loved ones, the community, and therapists who help them come to terms with what has happened.

1. Lapses that do not undermine the claim that one has an obligation

The slogan "ought implies can" is a bedrock commitment for many theorists and non-theorists alike. Those who accepts the slogan "ought implies can" might doubt that the first two agents have an obligation to help and the third an obligation not to harm. In order to respond to this concern, let us accept the slogan for argument's sake[11] and consider first whether there was a course of action that the agents were physically capable of executing and that would have resulted in the recipients being helped/not harmed. Given the following clarification of the details, the answer is yes. The neighbour could have taken the side streets from the start and not the freeway; the person on the pier, while not a first responder, is nevertheless sufficiently competent with a rope and could have thrown it on the first attempt where it would not have been swept away immediately by the currents (suppose that any place within some specified range to the right of the stranger will do). In the third case, there was a dry spot that was not in the pathway of the lamp where my friend could have stepped.

There are no other complicating features of the cases. For example, the person on the pier is not a doctor on her way to the emergency room in the wake of an earthquake. Doing something other than throwing the rope (or, in the first case, heading to the airport, and in the third case stepping where it is dry) is not simply less optimal, or less good, but wrong.[12] Upon seeing the stranger drowning, the person alone on the pier should not walk away, stare in despair or call 911. In each of the three cases, these agents correctly think there is a possibility of helping or avoiding harm and that other courses of action are not even remotely likely to have this outcome (e.g. there is no one the neighbour can call to go the airport on her behalf). In addition, what these agents must do is neither burdensome nor dangerous.

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

While these agents could physically perform an action that would have helped/avoided harm, could they have rationally perform that action? That is, was there a good course of reasoning available to these agents that would have led them to help/not harm? Their lapse in recalling the relevant fact(s) (e.g. the freeway construction) blocked off the relevant line of reasoning. They could not reason to the conclusion to perform the relevant action (e.g. take the side streets from the get go) given their lapse; their deliberations were missing the crucial proposition. If they could not have come to the conclusion to do that which would have in fact been helpful/not harmful, then, it might seem, we cannot saddle them with an obligation to help/not harm whose discharge relies on their being able to come to this conclusion.

Rational abilities do seem just as relevant as physical abilities. So, if one is obligated to help, then it is reasonable to require the agent to reach the decision to do what is helpful.[13] This in turn implies that it is reasonable to require the agent to bring the relevant knowledge to bear as she reasons.

Is it reasonable or unreasonable to require, for example, the neighbour to bring her knowledge of the freeway construction to bear on her deliberations about what to do? She has no habitual way of getting to the airport and or phobias that play into her choice of route. She has remembered construction alerts in the past. And perhaps she will not ever forget to consider whether the freeway is under construction again. She did not forget because she was distracted. It is not that she could not have remembered, but rather that she did not. She simply forgot. Would a reasonable person bring the relevant knowledge to bear in a situation like this? The fact is that reasonable people sometimes do and sometimes do not. That is, reasonable people too sometimes suffer from simple lapses. The same goes for my friend carrying my child.

The case of the person on the pier throwing the rope might seem less straightforward. Perhaps we cannot help but suppose the rope thrower is panicked and that is why she does not bring her knowledge of the currents to bear on her first attempt. But not every agent would panic. This one did not. She has a pretty cool head. She is not overwhelmed by the emergency circumstances, she is not slow thinking, and she is familiar with currents and how to counteract their force. She is not confused or overly absorbed in some other project. And she has no desire, conscious or unconscious, to avoid helping. Nevertheless, the relevance of the currents did not occur to her until after the first throw.

In all these cases, the agents' memory lapses are simple or basic: there is no further causal explanation that will make a difference to our moral evaluation of that lapse.[14] It is not that they could not remember but rather than they did not.[15] Most of us have, on some occasion, forgotten, not recalled, or not picked up on a cue. When the lapses are rare, backup plans – the note on the fridge, the string around the finger and so on – are not required. Only in cases where lapses are a matter of course (e.g. for a person with a full schedule, or if someone is forgetful) is some form of backup plan is called for. But even if one thought that we are always required to form a backup plan so to remember, notice that backup plans

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

too can be foiled by simple lapses, such as when an agent notes in her daily planner to avoid the route under construction, but then simply forgets to check her planner.

In cases where a lapse is not simple or basic, whether it undermines the claim that one has an obligation depends on the explanation for the lapse. If an agent fails to remember that the road is under construction because she is preoccupied with her garden project or because she is a forgetful person and did not take any steps to guard against forgetting, then she is nevertheless obligated to pick her neighbour up. However, if one fails to remember a crucial fact due to a sudden and unusual migraine, one does not have an obligation. It is unreasonable to expect an agent to overcome this sort of pain and to function as a normal adult.[16] In cases where a lapse is due to panic, much depends on whether, for example, the agent has special training in how to deal with panic. In some cases, we do require that even ordinary people overcome their panic or anxiety. But we need not decide here whether and when panic rebuts the claim that an agent has an obligation to act, since none of the cases described involve panic. In all three cases the agents could have remembered but did not due to simple lapses.

2. Challenges to the dichotomy

From this point forward, I focus mainly on the airport example, referring to the neighbour who made the promise as F. If one agrees that F has an obligation to help, it may be tempting to conclude she violates this obligation. In Section 2.1, I describe the challenges such a position must face up to. In response to those challenges, or even independently of them, one might instead hold that F only has an obligation to try to help and since she does try, she satisfies her obligation. In Section 2.2, I describe the challenges facing this view.

2.1 Challenges to the claim that F violates her obligation

In this section, I consider two different arguments concluding that F violated her obligation to help. The first, perhaps convinced by the earlier discussion that F's obligation to help is not rebutted by her simple lapse, holds, in addition, that she is at fault for not bringing her knowledge of the freeway construction to bear on her reasoning, and thus she violates her obligation to help.

(1) F does not satisfy her obligation to help her neighbour because she does not bring her knowledge of the freeway construction to bear in her deliberations about what to do.
(2) F is at fault for not bringing her knowledge to bear in her deliberations.
(3) If one does not satisfy an obligation to act due to a fault of one's own, then one violates that obligation.
(4) So, F violates her obligation to help.

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

There are two ways to understand "fault." According to the first interpretation, the phrase indicates blameworthiness. Given this interpretation, I think the third premise is extremely plausible.[17] Take, for example, a lifeguard who does not succeed in helping someone who is drowning because she reacted too slowly, and she reacted too slowly because she freely decided to drink while on duty. She is blameworthy for her slow reaction time, and this does seem to explain why we evaluate her action as a violation of her obligation to help.

However, given this sense of "fault," is the second premise true? That is, is F's mental lapse blameworthy?[18] Consider some lapses that clearly are blameworthy. Imagine a greedy day trader who is so focused on making her next thousand that she fails to notice her office mate having a heart attack. Or suppose an agent is in a hurry to get to her movie on time and so carelessly throws a rope to someone drowning in the water. In each of these cases, the agents have an obligation to help and they fail to do so through a fault of their own. Although the explanation for the lapses in these cases varies, at bottom each lapse reflects or results from an inappropriate orientation or disregard, either temporary or of a more permanent nature, toward those in need of help. The agents care more for themselves or their projects relative to others than they should. However, F's epistemic lapse manifests and results from no such inappropriate orientation towards her neighbour.[19] F is wholly aimed at helping her neighbour.

F's epistemic lapse also manifests and results from no other blameworthy condition or aspect, such as unreasonable overconfidence. For example, suppose an agent is overconfident about her memory abilities and so does not write down crucial instructions. As a result she takes the incorrect steps in a project. Or consider the stranger described earlier who comes upon a drowning person, but now imagine that the stranger is overconfident about her rope throwing abilities. Rather than throw the rope, she decides to lasso the person. She is terrible at lassoing, and has sufficient evidence of this from her terrible performances at the past few rodeos. But she nevertheless unreasonably believes she can succeed in lassoing. Such an epistemic mistake is blameworthy. The trouble in these cases is not an inappropriate orientation towards others but rather an inappropriate (blameworthy) conception of oneself.

What about a driver who forgets to look before she turns? Looking where one is driving is one of the central norms that apply to the activity, barring special circumstances such as being shot at from above by a helicopter. It is so basic to the activity that to forget to look is like saying one forgot one was driving, and that is not the sort of thing one can simply forget. People of course do forget that they are driving (they are surprised as they pull up to their driveway), they forget to look where they are going, and they forget to keep their car in its lane. The explanation for these lapses is either blameworthy or not. Examples of the former are allowing one's mind to wander in a way that reflects a lack of sufficient regard for others on the road (including pedestrians) and overconfidence in one's ability to drive when tried or distracted or texting. An example of the latter is a very

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

unusual stroke that allowed one to somehow make it home but without much awareness of how one was doing it.

F's epistemic lapse is clearly different from all of the various types of above-described blameworthy lapses. It does not share any of the features that would make a lapse blameworthy.[20] F is not at fault, in the sense of to blame, for her epistemic lapse. If this is correct, then we have the following interesting result with respect to epistemic lapses: an agent meets the relevant conditions for epistemic responsibility (e.g. rational capacity, etc.), and yet her failure to do what she epistemically should does not entail that she is blameworthy for that epistemic lapse.[21] F had an obligation to bring her knowledge of the freeway construction to bear on what to do and she did not, but she is not therefore to blame for this epistemic lapse.

There is another way to interpret the word "fault" in the above-mentioned argument, namely as indicating a malfunction – a malfunction attributable to her, as opposed to the rope (its coming unwound, for example). It is true that F's recall did not operate as it should have, and this accounts for why she does not succeed in helping her neighbour.

Given this new interpretation of "fault," the third premise loses plausibility. Mere malfunction is insufficient grounds for the violation of an obligation. In fact, in some cases malfunction undermines the claim of an obligation. When a lifeguard does not help a drowning person due to a sudden migraine or a blood clot, she has not violated her obligation to help but rather no longer has any such obligation.

Those who wish to claim that F violated her obligation to help her neighbour can turn their sights to a second argument, which takes the obligation to help at face value, namely as requiring the result that the relevant person be helped.

(1) If one's action does not result in someone else's being helped, then one has not helped.
(2) If one has not helped the person one is obligated to help, then one has violated this obligation.[22]
(3) F's action does not result in her neighbour being helped and F has an obligation to help her neighbour.
(4) So, F has not helped the person she is obligated to help.
(5) So, F has violated her obligation to help.

One motivation for the second premise, the centrepiece of the argument, is the desire to take seriously the results of action. The results of an action surely do matter and in at least two ways. First, the results can make a difference to what action one has performed. Whether one's action is a helping or merely an attempted helping (or a killing vs. close call) depends, at least partly, on whether the person in need is helped (or whether anyone died). Second, the results of an action can make a difference to the evaluative distinctions we draw among actions. We would much rather be an agent who tried to help and succeeded than an agent like F, who tried to help but failed. Successful helpings are preferable

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

precisely because of the nature of their results – someone gets helped. Another motivation for the second premise might be the thought that agents like F should not get off the hook, morally speaking. F is morally responsible for the results of her action.

I agree with all of this. However, one can capture these various motivations behind the second premise by embracing the category of mere failure, as I will discuss in Section 3. Not only is there no need to accept the second premise but also there are certain challenges facing it.

How might such a view properly distinguish between F's action and V's action in the following case?[23]

> V has promised to pick her neighbour up at the airport. But she wants to watch the season finale of her favourite TV show and she knows that eventually her neighbour will make it home on her own.[24]

The results of V's and F's actions are the same – the neighbour waits for a long time and then takes an expensive taxi ride home. However, V's action (continuing to watch TV) is quite different from F's action (attempting to help the neighbour). More importantly, there is a difference as to which kinds of emotional reactions are justified. V, were she to have a moral moment of critical self-reflection, would be justified in feeling guilt, the neighbour resentment, and third parties indignation. But as I noted in the introduction, the emotional phenomenon we are presuming is that F is not to blame.[25]

There are a variety of next steps one could take, each incompatible with the others. (1) One might reject the presumption that F is not to blame and instead insist that both V and F are to blame, though V more so than F. (2) One might draw a distinction between a person being blameworthy and others being justified in blaming her, and claim that the acts of V and F are both equally blameworthy but that we are not justified in blaming F at all while we are justified in blaming V. (3) One might instead suggest that, though both V and F violate their obligation to help, we are neither justified in blaming F nor is F's act blameworthy. F's action is a case of blameless wrongdoing. I will address the first two points briefly and spend the rest of this section addressing the third.

Why might someone blame F (though less so than V)? Perhaps because they think F is responsible for her failure to pick her neighbour up at the airport. I agree that F is on the moral hook, but, as I will show in Section 3, we need not reach for blame as the only way to register responsibility. More importantly, there is not merely a difference in degree but a difference in kind between the actions of V and F. V's was the result of a blameworthy orientation towards her neighbour, while F's action was not. Using degrees of one and the same reactive attitude (blame) fails to reflect this difference in kind.

Suppose one agrees with the presumption with which the article began to at least this extent: blaming F to any degree is unjustified. If the justification for blaming is distinct from the justification for blameworthiness, then one could maintain that F's action is blameworthy (because F violated her obligation) but

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

that we lack the necessary standing to justifiably blame F. F's failure is all too familiar. We ourselves have messed up in just the same way and it would be hypocritical of us to blame F. "He who is without sin among you, let him be the first to throw a stone at her" (John 8:7). But notice that it is not only inappropriate for us to blame F but it also inappropriate for F to blame herself.[26] Guilt seems just as much out of order as does resentment. F did not act out of a blameworthy orientation towards her neighbour, while V did. Thus, both the blameworthiness of F's action and blaming in all its forms (guilt, resentment and indignation) does not seem justified in the case of F.[27]

Is F's action then a case of blameless wrongdoing, that is, a case in which she violates her obligation to act but blaming and blameworthiness are not justified? One way to explain F's blamelessness is that F has an excuse for her wrongdoing.[28] F and V both act wrongly, but F is fully excused due to her non-blameworthy mental lapse while V has no such excuse.

Consider these classic examples of blameless wrongdoing to determine whether F's action is sufficiently analogous to them.

> A person comes home at night and flips the light switch in the usual way as she has every night in the past. But on this night the flipping of the switch creates an electrical charge that flows next door and electrocutes her neighbour to death. There is no reason the agent should have suspected that her switch was faulty.
>
> An agent has drugs forced upon him, which cause him to intend to kill another, which he then succeeds in doing.
>
> A person is thrown off a cliff and kills another person on whom he lands.[29]

The line of reasoning which leads some to categorize these actions as blameless wrongdoing is as follows. All moral agents have an obligation not to harm or kill. For example, had the first agent come home and stabbed her neighbour to death, she would have acted wrongly. Since she killed her neighbour by flipping the switch, she violated this obligation. However, in this case blame would not be justified; her blameless ignorance about the dangerous light switch excuses her from any blame for killing her neighbour. Similarly for the other two agents.[30]

Even if one were to accept that the above-mentioned actions are instances of blameless wrongdoing,[31] notice that F's action is importantly different from these agents' actions. While it is not reasonable to require the agent to know that flipping the switch will kill her neighbour (and it is not reasonable to require the agent thrown from the cliff to land somewhere other than where he does, etc.), in the case of F it is reasonable to require her to bring her knowledge of the freeway construction to bear on what to do, to reach the conclusion to take the side streets to the airport, and to then execute that action.

Is F's action a better example of blameless wrongdoing than the purported examples? On the one hand, it seems to do some violence to ordinary language to call an action "wrong" while also maintaining that it is blameless. An action's being wrong seems to involve more than it being bad, bad for the victim, unfortunate, or regrettable. To call an action "wrong" naturally invokes the kind

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

of criticism that goes with blame.[32] On the other hand, we can use the words any way we like so long as we avoid confusion.

The more pressing problem is this. Recall the starting point of the article, namely that F is justified in having some sort of critical emotional response to herself due to how her action affected her neighbour. What is the thought content associated with F's feeling? If the constitutive thought of her emotional reaction were "I did the wrong thing" or "I violated my obligation to my neighbor," which according to the view under consideration would be correct for her to think, then her feeling would be no different from guilt (self-blame). And so she would be justified in blaming herself. And yet the view under consideration claims she is blameless and thus not justified in blaming herself.[33] And so those interested in labelling F's action as a violation must face up to the challenge of explaining how the feeling F is justified in having, which involves placing herself on the hook of moral responsibility, differs from that of self-blame.[34]

2.2 Challenges to the claim F satisfies her obligation

Either in response to considerations like those discussed thus far, or perhaps just as a direct response to cases like F's and even the light switch case, it is natural to feel pulled by the view that we should not expect too much of one another. We are only human. We are physically and rationally limited and we should not be saddled with obligations that are beyond our capacities. And so one might conclude that all we can really demand of one another is that we try our best; that is, we should do our part (by acting with the appropriate intentions, effort and care) and the rest is out of our hands (and perhaps in God's hands, or up to the world).[35] On such a view, obligations to X (where X might be helping, not killing, etc.) are not to be taken at face value, that is, as requiring that one X. Instead the obligation to X is typically reinterpreted as only:

An obligation to adopt the end of Xing[36]

An obligation to act with the intention of Xing

An obligation to try to X

Each interpretation makes transparent the satisfaction conditions of an obligation to X. Depending on how each is interpreted, they are not obviously equally demanding,[37] but in any case F meets all these obligations (she adopts the end of helping, acts with the intention of helping, and tries to help), and so the possible differences among these interpretations will not affect the discussion in this article. Insofar as one is interested in maintaining that F neither failed nor violated her obligation to help, one might embrace the claim that F's obligation was merely to try to help (or one of the other above-mentioned formulations) and since she did so, she satisfied all the obligations she had.

In spite of the attractions of re-interpreting an obligation to act in these various ways, there are some challenges facing such a view. First, treating *every*

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

obligation to X as merely an obligation to try to X (with the appropriate level of effort and care) does not fit well with fact that we sometimes invoke the full-fledged obligation and in other cases we invoke merely the scaled back obligation (only the obligation to try).[38] If someone is in a position to easily help and there are not any reasons against helping (e.g. opening the door for someone in the normal situation) we say that she is obligated to help, not merely to try to help. But if, for example, someone is blamelessly ignorant of the fact that the door is stuck, and reasonably believes that opening the door is an easy way to help another, then we say she has only the scaled back obligation to try or to intend to help.[39]

Suppose we accept that it is only in some, not every, circumstance that one is only obligated to try to X. Is F's situation like this? I do not think so. At the point when F first leaves the house, it is possible for her to pick up her neighbour on time. Moreover, it is reasonable to require F to know that the only way to get to the airport on time that day is via side streets. The claim that F has an obligation to help her neighbour does not place an unreasonable or impossible burden on F.[40]

A second challenge facing the view is its ability to properly distinguish between the actions of F and S in the following case.

> S promises to pick her neighbour up at the airport. Remembering that freeway construction has begun, she leaves early and takes side streets to the airport. She arrives on time and brings her neighbour safely home.

Both F and S set out to help their neighbour, but S, unlike F, takes into account the freeway construction and so succeeds in helping her neighbour. Since both act with the appropriate intention, the view must treat both as having satisfied their obligations. The question is how to account for the result of S's action. The neighbour's getting picked up on time is not simply a bonus. There is a difference as to kinds of self and vicarious emotional reactions that are justified. F, unlike S, has something to feel bad about. This feeling is not mere regret, which is something that any bystander would be justified in feeling. F will be justified not only in feeling bad that she did not achieve her end, but also in feeling bad for failing her neighbour. Her feeling involves the thought that she is on a moral hook that mere bystanders are not. That this first-personal moral and emotional reaction is justified does not sit well with treating F's case as a success story.

One might resist the idea that F failed her neighbour because it seems that moral failure and success should be something under our control. The thought might go like this: there is much that happens in the world that is beyond our control, which is not the sphere for moral theorizing, though it is a sphere for mere regret.

My response is two-fold. First, F did have the relevant kind of control. She could have remembered the freeway construction, but she did not. Second, the sphere of morality and moral concerns does not neatly overlay with the sphere of control. For example, we are not always in control of whether we come to have

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

debt of gratitude to another person and yet in such cases it often (not always) is a moral failing not to be grateful in response to the one who helped us. Agent-regret is yet another moral phenomenon where a certain kind of control is absent and yet we end up standing in a new and problematic moral relation to another, where repair is called for, when and if possible.[41] So there are indeed important areas of morality where we lack control and they deserve to be theorized about and to inform those parts of our moral theorizing that received the lion share of attention.

In light of this last comment about agent-regret, one might embrace it and respond to the challenge of accounting for the difference in emotional reactions in the cases of F and S by claiming that a more precise statement of what F is justified in feeling, which coheres with what has been said thus far, is agent-regret. This would allow one to make sense of the fact that no obligation has been violated while also maintaining the agent is nevertheless justified in feeling that she has failed her neighbour; she is morally responsible for what she has put her neighbour through.

But the phenomenon of agent-regret is a distinctive response to the unfortunate results of one's action in spite of the fact that there was no rational deliberative access to doing what would have been needed to avoid harm or to help. Williams introduced the example of the lorry driver (who could not know, through no fault of his own, that driving down this street at that time was going to lead to his killing a child) as a case in which an agent meets all the rational requirements of deliberation in order to criticize a Kantian account of justified action in terms of meeting rational requirements of deliberation. While the lorry driver could have physically done otherwise, and thus avoided killing the child, this option is not rationally deliberatively accessible. This is not what F's situation is like. F, unlike the lorry driver, does *not* meet all her deliberative requirements. She did not bring her knowledge of the freeway construction to bear on her reasoning but she could have and should have (Section 1).

One might try to insist that agent-regret covers two very different kinds of cases (F and the lorry driver). But the content of F and the lorry driver's emotional reaction is importantly, morally different. Just as there is a reason to distinguish agent-regret from guilt in light of the different thought content constitutive of, or associated with, these feelings (the former involves no thoughts about wrongdoing while the latter does), so too there is reason to distinguish agent-regret from the way in which F feels bad for failing her neighbour.[42] The content of F's thought is that she did not do what she should have, which is correct. The lorry driver has no such thought, which is also correct. This is a key moral difference in their emotional reactions and it should inform the kind of repair called for in the two cases.

This brings us to a related challenge regarding repair. The view that to satisfy one's obligation to help one need only act with the intention to help has trouble making sense of the need to make up for what one has done (when possible) in cases like that of F. To see that this is so, consider apologies. Apologies (often

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

expressed by the phrase "I'm sorry") have a wide array of uses, many of which have nothing to do with admitting responsibility for failing another. For example, sometimes apologies are used to offer sympathy, such as when one hears of the death of a friend's loved one. In other cases, an apology is used to indicate that the circumstances or results of one's action are unfortunate or regrettable (even when the action itself is justified), such as when someone goes without medicine due to scarcity. But in some cases apologies are also used to acknowledge that one has failed the other person (and even, in some situations, to thereby to make it up to the other person[43]).

Apologies, as used to acknowledge that one has failed another, not only make sense in cases where one has acted with an inappropriate intention, but also in contexts like that of F who did act with an appropriate intention. It is appropriate for F to apologize, but the challenge for the view under consideration is to explain why this is appropriate.[44] What, on this view, is she apologizing for? It would be inappropriate to accuse her of not trying hard enough. This is not the correct charge in the face of a one-off simple memory lapse, and so that is not what she should apologize for. According to the view, she is only obligated to try to keep her promise and she does try to keep her promise, so she satisfied her obligation with respect to her neighbour. Of course she might be frustrated that she did not succeed in keeping her promise, but that is no reason to apologize to her neighbour, since that was not owed to the neighbour. And yet it is clear that, at the very least, an apology is in order. She did not make it to the airport at the time she promised and it was reasonable to expect to her to do so.

One might think that while all the intentions that F has are appropriate, she is nevertheless missing an intention she should have, and this explains why she should apologize. She should not only adopt and pursue the end of picking up her neighbour, but she should also adopt and pursue the mediate end of taking side streets rather than the freeway under construction. If the appropriateness of an agent's mediate ends is judged relative to what it is reasonable to require an agent to believe, rather than relative to what an agent does believe, then F is missing an appropriate intention and thus does not try to do what she should try to do (take the side streets). However, the problem with this line of thought is that, recall, the appeal to an "intention only" view of the satisfaction of obligations was meant to *defend* the claim that F satisfies her obligation to help her neighbour.

Here then is the second challenge presented in full for those who wish to hold that F neither merely failed nor violated but instead satisfied her obligation by appealing to the claim that all we are obligated to do is act with the appropriate intentions. If one accepts that F acted with all and only the intentions she should have, then one cannot properly distinguish the actions of F and S, explain why F is justified in feeling bad for failing her neighbour, and make sense of the phenomenon of apologies. However, if, instead, one believes that F is missing an appropriate intention, then one can no longer claim that F satisfied her obligation.[45]

I would like to discuss one last challenge, which is more easily illustrated using the example of the person on the pier throwing the rope to a drowning

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

stranger. Suppose the rope thrower adopts the end of satisfying her obligations (this need not be her end, but suppose it were) and she believes the view under consideration, namely that what she is obligated to do is to try to help (or act with the intention of helping). So long as the stranger is in need of help, she will stay engaged with the situation and throw the rope again with the intention of helping.[46] However, she will have the wrong conception of her own activity. The first throw (at t1) cannot be thought of as an attempt, since with the first throw she met her obligation to try to help or to act with the intention to help. The use of the description "attempt" only makes sense when success is in doubt or one has failed. But there is no doubt that she will and has succeeded in trying at t1. Since her first action must be viewed as successful, her second throw cannot be conceived of as an attempt to succeed where the first attempt failed. That is, she cannot conceive of her second throw as "trying again" to satisfying the obligation she had at t1. The fact that the rope floated away from the drowning stranger, and the fact that he is still drowning, gives her new information and gives rise to a new obligation, namely to try to help or act with the intention of helping at t2. This is the same sort of information she could have gotten from watching someone else throw the rope without taking the currents into account. Her connection to the situation upon the second throw is, in this sense, fresh. Rather than conceive of herself as in the middle of something, that is, in the process of satisfying a single on-going obligation to help, which requires that the recipient be helped, she will conceive of herself as taking up something for the first time – namely, satisfying a new obligation of trying to help or acting with the intention of helping, which turns out to have the same content as the obligation she had at t1. Thus, the last challenge for the view that reinterprets an obligation to help as merely an obligation to try to help, or to act with the intention of helping, is to make sense of the phenomenon of "trying again" to satisfy one's obligation.

3. Benefits and implications of embracing the category of mere moral failure

A key advantage of giving up on the dichotomy of evaluation and accepting the category of mere moral failures is being able to properly account for why F is justified in feeling bad, but not justified in feeling guilt (self-blame), mere regret or agent-regret. Since F has not violated her obligation to help, she has not done anything wrong, and so guilt is not called for. At the same, time F had an obligation to help her neighbour, which requires that her neighbour be helped, and so she did not satisfy this obligation. That is the moral hook she is on. So, she has something to feel bad about that is quite different from what one feels bad about in the case of mere regret or agent-regret, where no such obligation is in place. F has reason not merely to see the trouble her neighbour has gone through as bad for the neighbour and that it would have been better for the neighbour if things had gone differently – the key constitutive thoughts of mere regret. And F has reason not merely to recognize that this bad outcome resulted from her action and that it would have been better had F acted differently (e.g. taken the

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

side streets from the start), which are the thoughts constitutive of agent-regret. F also has reason to think that she *should* have acted differently and that she failed her neighbour. Each of these four categories (being blameworthy, merely failworthy, agent-regretworthy and merely regretworthy) indicates a distinctive relation we have to our intentions, obligations, and the results of our or another's actions.

The feeling F is justified in having is a type of moral anger and disappointment with herself whose key constitutive thoughts are that she failed to satisfy her obligation and is on the moral hook.[47] What are we to make of the fact that we have no term (at least in English) for F's distinctive sort of feeling? We also did not have a name for agent-regret until Williams coined it as such, but that does not mean we were not familiar with the phenomenon or that it needed a name in order for it to pose a challenge to certain kinds of moral theorizing. As Aristotle pointed out, we do not have names for all the virtues and vices that pertain to emotions, and so it should not surprise us that the range and fine-grainedness of our emotional reactions can outstrip our current labels. The fact that we have lacked a name for mere failures, and that this category has been overlooked until now, might explain why some agents in situations like F's sometimes say they feel guilty – perhaps they reach for this term since it accurately captures their being and feeling responsible. However, since guilt also implies wrongdoing, they also thereby have an inaccurate construal of their action.

In addition to making sense of F's feeling of responsibility without being blameworthy, there are yet other benefits of accepting the category of mere moral failures. First, it allows for us to properly distinguish the moral status of F's action from the actions of S and V. Intuitively, F's action should not be lumped together with either. The difference is not one merely of degree but kind. S's action satisfies, V's action violates, and F's action merely fails with respect to the obligation to help. Second, it also allows us to make sense of the need to acknowledge the failure and attempt to repair it, of which the phenomenon of apologies is one example. F is apologizing for her failure to satisfy her obligation to her neighbour. In the rope throwing case it is not possible to make up for what one has done. But there are forms of symbolic repair that agents sometimes engage in: going to a temple and making an offering, or committing oneself to a new project in honour of the one who died, etc. Third, we can now make straightforward sense of the phenomenon of "trying again" in the case of the roper thrower. She has one on-going obligation to help, her first attempt failed to satisfy it, and so she tries again to satisfy it.

Broadening the basic categories for act evaluation to include mere moral failures has implications for what the necessary conditions are for both violation and satisfaction of an obligation to act. When it comes to violating an obligation to help or not harm, not helping or harming is insufficient for violation (as the difference between V and F makes clear). The agent's not helping must be due to a blameworthy orientation of the agent's will or due to some other blameworthy condition of the agent.[48] It is natural to focus on the fact that an agent did not help

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

and whether she could have helped as the necessary conditions of the violation of an obligation to help. Much of the literature has focused on whether and in what sense the ability to do otherwise is a necessary condition for violating an obligation.[49] However, I am highlighting this necessary condition for violating an obligation to act: a blameworthy orientation of the agent's will or other blameworthy condition of the agent (see Section 2.1).

Is action due to a blameworthy orientation or condition of the agent a sufficient condition for violating an obligation to act? Strawson thought that manifesting ill will or a disregarding will towards another was sufficient condition for resentment (blame). And justified resentment implies a violation has occurred. But it does not necessarily imply that a violation of an obligation to help or avoid harm has occurred. Consider these two cases:

> A jet skier sees a child drowning nearby. There is no one else nearby to help the child. The agent is practicing for an upcoming competition and does not like anything to interfere with her routine. So, she skis away from the child. However, in doing so, she unknowingly creates a wave that carries the child safely to shore.
>
> Alfred, seeing his wife's inheritance, gives his wife something to hasten her death but which in fact cures her of an ailment neither knew that she had.[50]

Both the jet skier and Alfred act in a way that manifests an inappropriate orientation towards others, and while they both no doubt violate some sort of obligation (an obligation to have, abide by and act on and out of an attitude of respect), it does not seem that they violate an obligation to help or to avoid harm. Violating an obligation to help, avoid harm, etc. requires more than merely having and acting out of ill will. So, there are at least two conditions are necessary for violating an obligation to act: (a) failure to do the act (b) due to blameworthy orientation or condition of the agent.[51]

A second implication of accepting the category of mere moral failures is that acting with the intention to help or not harm is insufficient to satisfy an obligation to help or not harm and that the relevant results are necessary. An obligation to help or not harm can be taken at face value as requiring a certain result.[52]

Is the intention relevant at all, or are certain consequences sufficient for satisfying an obligation to act? Looking back to the examples of jet skier and Alfred might help. Do they satisfy their obligation to help and not harm? If you think not, then you are committed to the view that merely doing that which helps or avoids harm is by itself insufficient to satisfy an obligation to help or not harm.[53]

Acting with the relevant intention and succeeding might also not be sufficient for satisfying an obligation to act. Consider whether the following agents satisfy their obligation to act.

> An agent intends to help and succeeds in helping but she mistakenly thinks that helping is merely optional and not obligatory.
>
> An agent intends to help her friend and succeeds in doing so, but mistakenly believes that it is obligatory to help all and only her friends.

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

> Upon hearing that Echinacea will cure her daughter (which is not true), a mother picks some from her garden and gives it to her daughter. Had she bought Echinacea from the market, the daughter would still be ill. However, due to unknown features of the soil in her backyard, the herb she picked does cure the daughter.

Full exploration of this issue must be left for another topic, but I am inclined to think that the agents have not satisfied their obligation to help in these cases and so further conditions beyond intention and bringing about the relevant consequences are necessary for satisfying an obligation to help.

What is clear thus far is that neither an agent's intentions alone are sufficient for satisfying an obligation to act nor do the results of an action alone settle whether one has violated an obligation to act. And so one should reject the "either or" form of the question, "Is it the results or the agent's intention that matters?"

One need not worry that requiring an agent to succeed in helping is overly burdensome. *If* agents have an obligation to help, then they must succeed in helping. But, *whether* agents have an obligation to help will depend, among other things, on what their other commitments are, as well as their physical and rational capacities.[54] This accommodates the fact that humans are physically and rationally limited beings. In fact, insofar as one is concerned to put obligations within the reach of human beings given their limitations, I do not see that replacing obligations to act with obligations to act with an appropriate intention will be of any help. Human beings are no more in a position to guarantee that they act on appropriate intentions then they are to guarantee the results of their actions. Both require, among other things, knowledge, discipline, will power, lack of the vices of character, and the cooperation (or non-interruption) of the world.[55]

Conclusion

I will conclude by indicating the range of actions that I believe belong to the category of mere failures. Thus far I have noted that mere failures arise both with respect to so-called perfect obligations (e.g. not to harm) and imperfect obligations (e.g. to help). Moreover, I have focused on mere failures that result from the agent's forgetting or not recalling a fact necessary for taking the correct means. For example, F does not recall the freeway construction. What if one forgets to do the very thing one has promised, such as return a borrowed book on a specified day? Returning the book on that day is not merely instrumental to satisfying her obligation but constitutive of the very act she is obligated to do. And yet it does seem the sort of thing one can simply forget.

Not only are there mere failures due to a lapse in thought,[56] one could apply the concept of simple or basic lapses to the realm of execution. For example, imagine that the rope thrower had remembered the currents but then bungled the throw. She has the ability to throw it in the correct location and she intends to do so, but she simply botches the execution of this physical manoeuvre. Suppose she is not clumsy or inattentive. Her lapse is simple. Had the rope thrower bungled her throw in this way, her second throw would have arrived too late. Simple lapses in execution are not uncommon even among world-class athletes.

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

A professional tennis player can be focused and determined but nonetheless make an unforced error. Sometimes a football receiver simply drops the ball. Or an agent trips and breaks the watch she was attempting to return.

Simple lapses in execution are, in one sense, inevitable. Not even the greatest tennis player in the world will hit an ace with every serve. Likewise, not even the greatest lifeguard will always successfully throw the rope to those who are drowning. But the mere fact that human beings are finite in this respect, that is, that we are subject to such simple lapses, does not itself undermine the claim that one has an obligation to help in a particular situation. Of crucial importance is whether it is reasonable to require an agent to do what will be helpful in circumstances she faces. I think it can be reasonable to make such demands on agents even in the face of simple lapses in execution (Section 1).[57]

In some cases, it will be hard to know what sort of situation we find ourselves in. That is, it may be difficult to determine whether there really is an obligation to help and, if so, whether one has violated it or merely failed with respect to it. Moreover, hindsight does not always clarify the matter. In some cases we may never discover or settle to our satisfaction whether we unreasonably overlooked some crucial fact.

What is clear is that good intentions, care and effort are not sufficient for satisfying an obligation. It is possible to adopt the proper end, sincerely pursue it, take care while doing so, and yet merely fail to satisfy the obligation. I think this is one of the unsettling, but perhaps distinctive features of the human condition: we can be put in a position where effort, care and good intentions are insufficient to compensate for our human limitations.

So many of our discussions are around who is to blame and for what. What this article shows is that we need to widen our moral focus beyond that of blame. Even if one settles that one is not to blame, one has not necessarily finished settling the moral issue at hand. An accurate appreciation of what we have done (or what has been done to us) involves acknowledging that there is more than one way of failing another person. In cases of mere moral failure one did not do what one should have, yet without violating that moral obligation. The next step is to examine how this category of "mere moral failures" might impact legal and social responses to a wide range of errors (e.g. liability for medical mistakes) including whether mere failure calls for a distinctive kind of repair.

Acknowledgements

I want to thank Stavroula Glezakos, Barbara Herman, Richard Kraut, Dana Nelkin, the blind referees at the Canadian Journal of Philosophy, and audience members who gave critical feedback on earlier drafts presented at University of Illinois, Urbana-Champaign, University of California, Riverside, California State University, Northridge, University of California, Irvine, University of California, Santa Barbara, Bay Area Forum for Ethics and Law, Ninth Southern California Philosophy Conference, First Annual Society for Ethical Theory and Political Philosophy at Northwestern University and the Work in Progress Group at the Claremont Colleges. I also want to thank David Beglin for proofreading assistance.

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

Notes

1. Although I am treating "obligation" as interchangeable with the terms "requirement," "ought," "should" and "supposed to," an alternative approach can be found in Simmons (1979, chapter 1).
2. There are, of course, many categories of evaluation, such as "permissible" and "supererogatory," but these forms of evaluation are not at issue in this article.
3. In this article I focus on moral obligations and aim to establish the category of mere moral failure, though the term "moral" is often left implicit. However, I also believe that one can merely fail with respect to non-moral obligations.
4. This is not a case involving an implicit promise (and hence a version of the first example), but rather a situation where an agent is set on avoiding harm to another, and thus focused on one feature necessary for doing so, but who forgets to consider a second feature necessary for avoiding harm.
5. In the first section I provide further details to persuade the reader that aiding in the second case is obligatory, not supererogatory. The reader may, in the end, find the other cases more convincing. Since no one case may speak equally strongly to all audience members, and any one case suffices for my purposes, I present a range of cases so that readers can focus on the one they find most compelling.
6. To say that an emotional reaction is justified is different from claiming it is typical, understandable, useful, or required. It might be typical for LA drivers to become enraged when cut off in traffic, understandable for parents to become extremely fearful when their child faces a minor danger, and useful to be irate with the customer service representative since this makes it more likely one will get reimbursed. However, if a feeling is justified, then the content of the thought that constitutes or is associated with the emotion is correct (as well as the emotion's duration, intensity and so on). See, for example, Aristotle's *Nicomachean Ethics*, Books III–V, D'Arms and Jacobson (2000), and Tannenbaum (2007). Fear, for example, involves the thought that the object is dangerous and this emotion is justified only if the object is in fact dangerous (there are complications here that I do not have space to address such as when one's belief that something is dangerous reasonable but false). Lastly, note that a feeling can be justified without being required: being nervous may be justified when performing in the Olympics, but one makes no sort of mistake if one instead has only feelings eagerness or competitiveness.
7. For some readers, the key aspect of this article will be (a) and (b), namely that there is a reactive attitude reflecting the agent's responsibility for morally failing that falls between blame and mere regret, and, as I later show, is importantly different from agent-regret.
8. In fact, the landscape is even more complicated since in some cases an obligation is taken care of by another before one has chance, is transferred to another, etc. But I am setting such cases aside since they are not relevant to the issue at hand.
9. Aristotle relied on emotional reactions to inform, among other things, his account of the counter-voluntary, the virtues and the difference between the enkratic and virtuous agent. Strawson (1962) highlighted emotional reactions such as resentment and focused on the sort of considerations that mitigate or remove it. He then used this as a basis for critiquing certain theories about moral responsibility and free will. Williams (1981) criticized a Kantian account of justified action by appeal to an emotional reaction he dubbed "agent-regret." Many others (e.g. Stocker [1976]), have also appealed to emotional reactions to critique and broaden moral theorizing.
10. This article will not address the question of what sort of repair may be called for in cases of mere moral failure. Rather, its aim is to establish that there is such a moral category.
11. There is a large literature on whether to accept the slogan and how to interpret it. For a recent discussion, see Scanlon's (2008, 45–52) objection to Judith Jarvis

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

Thomson's use of an "objective ought." Nothing in this article hangs on accepting or rejecting the slogan. Here I am only showing that accepting the slogan, at least under one reasonable interpretation, is consistent with claiming that the agents described earlier have an obligation to help/not harm.

12. While the drowning person would be no worse off were the agent on the pier to do nothing, this does not lessen responsibility of the person on the pier. This is akin to the way in which doctors have obligations to help their patients even when the patients would have been no worse off were the doctors to do nothing. Insofar as one thinks that some cases of the imperfect duty to aid can be obligatory, (e.g. Singer's [1972] example of a drowning child), one should fashion such a case in one's mind.
13. Some epistemic obligations have theoretical, others practical sources, and within the practical, both moral (e.g. generated by making a promise) and nonmoral (e.g. generated by a desire for contentment) sources.
14. What makes a lapse simple has nothing to do with whether it is extraordinary, uncommon, or unexpected, but rather with the explanation of the lapse – whether the causal explanation is such that it makes a difference to our moral evaluation. For further examples of such lapses, see M. Smith (2004). Not only can lapses be simple or basic but so too can mistakes. See Austin's (1979, 200, note 1) discussion of mistaking the hot for the cold water tap. In Section 2.1 I discuss whether simple lapses ground blame or serve as an excuse from blame.
15. One might have this worry: these agents could not have recalled through an act of will, and thus they *cannot* recall. This would then lead us to doubt that they have an obligation to help or avoid harm. Is voluntary control, however indirect, as opposed to mere reason-responsiveness, a necessary condition for the ability to recall? See, for example, Fischer et al. (2007) and Hieronymi (2009a, 2009b). For purposes of this section I am assuming that there are situations in which an agent could have remembered, but did not, while leaving open what the correct account of this metaphysical possibility is.
16. It is an interesting question on what basis we categorize some situations as those where one could have remembered but did not as opposed to those where one could not have remembered. In both cases there are events in the brain and questions of freedom and its connection to agency arise. As I noted earlier, I am assuming that there are situations in which an agent could have remembered but did not, without providing an account of this type of metaphysical possibility.
17. Thomson (1989) endorses a similar claim.
18. Later I will consider whether F's failure to pick up her neighbour is blameworthy.
19. And so F's epistemic lapse does not reflect an objectionable evaluative judgment or commitment. For a discussion of epistemic lapses that do reflect objectionable evaluative judgments or commitments see A. Smith (2005).
20. Similarly, when my friend bangs my child's head into a lamp because he forgets to look up, his lapse is not due to a blameworthy orientation or any other blameworthy condition.
21. Dahl (1967), among others, claims that not doing what one should is sufficient for being blameworthy. F's case challenges Dahl's claim as applied to epistemic obligations (later I show that F's case challenges Dahl's claim as applied to obligations to act). F has a morally derived epistemic obligation (obligation to know certain facts and to bring that knowledge to bear in the right way in her reasoning). She does not meet this epistemic obligation, she is not blameworthy for it, and she is not excused from it, all of which indicates that she neither violates nor satisfies this epistemic obligation.
22. Utilitarians who allow talk of obligations might accept such a claim. And some Deontologists hold a modified version of this claim. For example, Thomson (1991) implicitly accepts this view: if an agent has an obligation to help and this obligation

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

is not overridden, cancelled or otherwise neutralized, then not helping violates this obligation. A contractualist such as Scanlon is also likely to accept it: "What makes an action wrong is the consideration or considerations that count decisively against it," (2008, 23) and I think he would agree that there are considerations that count decisively against F taking the freeway to pick up her neighbour from the airport.

23. Perhaps this is a good place to reveal my naming scheme. V is so-called because she violates her obligation, whereas F merely fails with respect to her obligation. S, who I will discuss later, is so-called because she satisfies her obligation.
24. Or consider this version: V prefers to take the side streets for the scenery, and she knows this entails leaving the house early. But she does not leave the house early because she wants to watch the end of her TV show.
25. However, there is another set of reactive attitudes that are justified, and they involve the thought that F is on the moral hook and has failed her neighbour (see Section 3).
26. An additional concern is whether my own past behaviour is the right kind of reason for refraining from blaming another. See Hieronymi (2004).
27. I take no stand on whether blameworthiness and blaming come apart, but show that even if they do, it would not help us properly distinguish V and F's actions.
28. For example, according to Rosen "we are hardly ever culpable for the wrong we do" (2002, 62). "If [one] has been neither negligent nor reckless in the management of [one's] opinion – then [one's] ignorance is blameless and so is the act done from ignorance" (63). Sher puts forward this condition for blameless wrongdoing: "an agent is not responsible for an act of whose wrongness or foolishness he is unaware unless (1) his failure to recognize the act as wrong or foolish falls short of satisfying some applicable standard, and (2) his failure to satisfy that standard is in turn accounted for by some combination of his constitutive attitudes and traits" (2009, 21, see also 87–90). F's action is not "explicable in terms of the interactions among the innumerable desires, beliefs, attitudes, and dispositions, many of them unconscious, that together make him the individual he is" (21). If Sher is using "responsible" to mean "blameworthy," then Sher would characterize F's action as blameless wrongdoing. If, instead, he sees these as conditions for responsibility, then, contrary to Sher, I maintain that F is nevertheless morally responsible and that certain negative moral reactive attitudes (but not blame) are justified – more on this in the next section.

    There are two other ways one might pull apart blame and wrongdoing. One could appeal to Scanlon's account in which wrongness is determined by whether one acted contrary to principles that could not be reasonably rejected whereas "the blameworthiness of an action depends, in ways that wrongness generally does not, on the reasons for which a person acted and the conditions under which he or she did so" (2008, 124) My response to this view is discussed in the second part of endnote 34. Alternatively, one might loosen the connection between the evaluation of an act and the evaluation of the agent who did the act, and claim that while F's action is wrong, F is not to blame, since blame-judgments are a matter of an agent's character. Thomson (1991) and Ross (1939) adopt a view like this. However, I do not think this strategy will be of much help since one should judge not only V and F differently but also one should judge their actions differently. For an additional criticism see Scanlon (2008, 127).
29. The examples are drawn from Thomson (1990, 229 and 1991).
30. It is unclear whether these last two agents really act, but, according to Thomson, "agency is no more required for violating a right than fault is" (1991, 302).
31. Those who accept "ought implies can" will reject the claim that these agents have violated their obligation, and hence will also reject categorizing their actions as blameless wrongdoing. If an agent is obligated to act, then, they will claim, she must be able to reason to the conclusion to do the act (where this includes omissions).

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

If one applies this lesson to the obligation not to harm, then the harms one is obligated to avoid are limited to those it is reasonable to require an agent to know about and so reach a conclusion to avoid. It is not reasonable to demand, for example, that the agent know that flipping the switch will kill her neighbour. And so, holders of "ought implies can" would maintain that the harm she inflicts is not the harm that she has an obligation to avoid.

32. Blame certainly implies wrongdoing, but if wrongdoing also defeasibly implies blame, then one might accept this defeasible biconditional: an action is wrong if and only if it is blameworthy. Accepting this biconditional is consistent with maintaining that wrongdoing is the more fundamental concept or explanatory notion.
33. Thus, the following way of carving up the conceptual terrain is problematic: when one has an obligation to act, one either satisfies or violates that obligation; if the latter, then either blame or some other reactive attitude is justified. The problem is that blame and this other reactive attitude are indistinguishable since they have the same constitutive thought: "the act was wrong." Some take guilt to involve not only the thought that one has acted wrongly and a feeling (feeling bad about what one has done), but also the thought that one deserves to feel bad. And so one might try to distinguish F's feeling from V's by claiming that F does not have this extra thought, while V does – this is why what V feels is guilt, while what F feels is not guilt. This is a controversial interpretation of guilt and it is not clear we should accept it. In any case, an explanation is owed as to why it would be inappropriate for F to have the thought that she deserves to feel bad for what she has done.
34. The same challenge applies to those who might describe F's action as objectively wrong but subjectively right (see Graham [2010]) or as "justified wrongdoing" (see Buss [1997]). Note that, in any case, F's action does not seem to be an example of "justified wrongdoing." Buss' case and F's case have this in common: the agents are blameless and yet accountable. But her agent is blameless because justified, whereas F is not justified. F can be reasonably expected to know both what would help and that F has an obligation to help, whereas in Buss' case we "suspect he was justifiably ignorant" (338).

    The challenge facing a view like Scanlon's is different. Blame, according to Scanlon, is not tied per se to wrongdoing. The latter is a matter of doing an action where considerations count decisively against it and not a matter of the agent's reasons for action (2008, 23, 124), whereas blame is about the "meaning of the action," which is determined by whether the action shows something about the agent's attitude towards others that impairs the relations that others can have with that agent (128, 141). Applying Scanlon's view, only V, not F, can justifiably blame herself, since only V has a problematic attitude towards others. Moreover, since the meaning of an action depends on the agent's attitudes or reasons for action (55), then it would seem that F's action does not have any special meaning for her relationship with her neighbour. And yet the moral relationship between F and her neighbour is impaired as is clear from the fact that F is justified in having some sort of negative self-reactive attitude and for initiating repair (see Section 2.2). This impairment is not because of F's attitude but because of what F did. Moreover, how is F to think about what she did, except, on Scanlon's view, as wrongdoing? In which case wrongdoing itself does have moral meaning, or, alternatively put, the meaning of an action is not exhausted by the agent's attitude, as Scanlon claims, but can be solely a matter of whether she did what she was not supposed to do.
35. The Stoics, Kant and Kantians are often interpreted as holding such a view. See, for example Herman's claim that "good intentions plus adequate care are enough (whatever the outcome, the agent has done what she ought)" (1993a, 98). Some virtue theorists too hold such a view. Foot claims, "it is not, of course, necessary that

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

charitable actions should *succeed* in helping others" (1978, 169, note 12). The view can be found beyond the two camps just mentioned. For example, according to Prichard, "an obligation must be an obligation, not to *do* something, but to perform an activity of a totally different kind, that of setting or exerting ourselves to do something" (1949, 35) and Rabbi Tarfon claims, "It is not your duty to complete the task, but you are not free to desist from it" (1962, 2:21).

36. In what follows I assume that adopting and pursuing the end of helping is appropriate. It may be that one must also intend to do one's duty, or to act well, while seeing helping as instrumental to, an instantiation of, or constitutive of doing one's duty or acting well. The arguments in this article do not hang on the resolution of this issue. Also, when I say that an agent acted with the intention of helping, I mean this to be different from, and inconsistent with, for example, acting with the intention of helping so as to later inflict harm (e.g. reviving someone in order to then torture him).
37. One can adopt the end of helping without ever pursuing this end (perhaps because the pursuit of this end just so happens to always be overridden by other ends whenever the occasion to help arises). And even if adopting an end entails taking the means in a certain situation, it is yet a stronger claim to say that the means must be pursued with this intention. Moreover, perhaps trying to help requires very little effort whereas acting with the intention of helping requires more effort (e.g. the expression of trying seems applicable for a child who just dips his toe into the cold water. "I tried," he can truthfully claim).
38. An obligation merely to try to help is a scaled back in comparison to an obligation to help since much less is required to satisfy it compared to the obligation to help
39. It is an interesting question when one has an obligation to try to help, for example, rather than an obligation to help. We certainly would criticize a person for not trying if she reasonably but falsely believed there was a way to easily help another. Perhaps this reveals two conditions of when we have scaled back obligations: (i) it is through no fault of the agent's own that Xing is impossible and (ii) the agent would be morally criticizable if she did not try to X. An alternative view to consider: one always has an obligation to try and sometimes also has an obligation to help. Though perhaps we should reject this view for the same reason we reject this view: one always seems to see and sometimes also sees.
40. Similarly, when the person on the pier first comes across the drowning stranger, it is possible for her to save him, and it is reasonable to require her to know that what would best help is throwing the rope where it will not be swept away immediately by the currents. And in the case of the friend carrying my child it is possible to step, and reasonable to require him to know to step, where it is dry and there is no low hanging lap.
41. See Williams (1981) and Tannenbaum (2007).
42. Wolf (2001) considers two equally negligent acts only one of which causes death to another. She suggests that it is a virtue to take responsibility for the consequences of one's action (13) and that the negligent killer should blame himself, which is determined not merely by whether the agent is blameworthy but also by whether agent-regret would be appropriate. In this article I am developing an account of responsibility that does not appeal to blame or agent-regret (unlike Wolf's account of the negligent killer). F's case is importantly different from those considered by Wolf, since F is not negligent (see Section 2.1), and hence not blameworthy, and, as I have just explained, agent-regret is also not an appropriate response for F to have towards her action.
43. Granting that an apology can repair less than successful activity does not commit one to a "balance sheet" view of action. One cannot do wrong with the intention of later apologizing, as if this would make everything come out "even." It is another question – not one addressed in this article – whether and when an apology itself can or should be enough to make up for or repair what one has done.

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

44. Herman (1993a, 97–98) suggests that agents motivated by duty will adopt maxims of response in such a situation, but it is unclear why doing so is a form of repair if the agent has previously acted on well-formed maxims and if acting on such maxims is sufficient for satisfying one's obligation.
45. If F is missing an appropriate intention, is she therefore, on this "intention only" view of obligations to act, blameworthy? That has traditionally been the response. But, given the emotional phenomena that are our starting point, F is not blameworthy. So one who holds the "intention only" view now must join with me in rejecting the dichotomy (violation vs. satisfaction) of the basic categories of evaluation and accept the category of mere failure. In addition, there is still room for disagreement about this: what obligation did F fail – merely an obligation to act with the appropriate intention or, what I suggest in Section 3, an obligation to help? The last challenge considered in this section helps push us yet farther away from an "intention only" view of obligation.
46. See Herman (1993b, 28–29, 1993a, 98).
47. According to Sher there are four different contexts in which we attribute responsibility. We do so whenever we (1) blame an agent for acting wrongly; (2) praise an agent for exceeding his duty or doing the right thing under difficult circumstances; (3) criticize an agent for acting foolishly or self- destructively; or (4) give an agent credit for his success in achieving some aim" (2009, 14). If Sher thinks that these are the only categories of responsible action, then our views conflict. As F's action demonstrates, there are actions that are not wrong, not blameworthy and neither foolish nor self-destructive, but for which the agent is nevertheless morally responsible.
48. Thomson (1991) claims that self-defence is justified in order to defend one's rights from violation and that innocent threats (that kill) do violate one's right not be harmed or killed. However, an innocent threat does not threaten due to a blameworthy orientation or condition. And so, were an innocent threat to harm, it would not violate an obligation not to harm. If violation of this obligation is necessary for a rights violation, and a rights violation is necessary for justified self-defence, then an alternative justification must be offered for self-defence against innocent threats.
49. See, for example Fischer (1986), Watson (2003) and Sher (2006). An exception is Strawson (1962).
50. Based on an example from Thomson (1991, 293).
51. Is the combination of failing to help or avoid harm and a blameworthy orientation or condition of the agent (where the latter explains the former) a sufficient condition for the agent's violating an obligation to help or avoid harm? This will depend on, among other things, whether and in what sense the ability to do otherwise is a necessary condition for violating these obligations, which I am not addressing in this article.
52. That is, *if* one has an obligation to help or not harm, then one's action must result in the relevant person being helped or not harmed. But if, for example, an agent reasonably but falsely believes road X is the only way to the airport (perhaps an otherwise reliable radio announcer mistakenly directs listeners to road X, which is in fact blocked), then she only has an obligation to *try* to pick her neighbour up from the airport, not an obligation to pick her up (see endnote 39). And if someone beats her to the punch and picks up her neighbour first, then she has no obligation to help (see endnote 54) or to try to help.
53. Notice that the jet skier, Alfred, and F fail to satisfy their obligations. However, the jet skier and Alfred violate at least one obligation (of attitude) while F does not.
54. Moreover, the existence of an obligation also depends, in very particular ways, on the facts of the situation. If sharks suddenly appeared on the scene in the midst of hauling

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

the drowning stranger to safety, and devoured the stranger in mere seconds, then the agent's attempt to save would not have failed her obligation. In such a case the obligation to help dissolves. Likewise, if someone had jumped in the water and saved the drowning stranger before the person on the pier had a chance to throw the rope, the rope thrower's obligation to help would have dissolved.

55. See Nagel (1979).
56. I have not by any means discussed the full range of lapses in thought that can lead a person to fail to satisfy her obligation. For example, I have not discussed errors in reasoning that lead an agent to undertake an unsuccessful course of action. Consider, for example, an engineer who makes a simple multiplication error while calculating the force necessary to stop an empty trolley from barrelling into people trapped on the tracks. As a result, his use of force was inadequate to stop the train.
57. While lapses are inevitable (they will happen), it does not follow that they are unavoidable (they have to happen; they cannot be avoided). And so one still has an obligation to help or not harm, as opposed to an obligation to help most of the time or merely to try to help. Consider a different sort of case, namely weakness of will. It too is inevitable but not unavoidable. It can be true that I will succumb to temptation, while also being true that I can and should resist temptation.

Notes on contributor

Julie Tannenbaum is an associate professor in the philosophy department at Pomona College. She works on philosophical issues that arise in moral psychology, ethical theory, metaethics and bioethics.

References

Austin, J. L. 1979. "A Plea for Excuses." In *Philosophical Papers*, 175–204. New York: Oxford University Press.

Buss, Sarah. 1997. "Justified Wrongdoing." *Nous* 31: 337–369.

Dahl, Norman. 1967. "'Ought' and Blameworthiness." *Journal of Philosophy* 64: 418–428.

D'Arms, Justin, and Daniel Jacobson. 2000. "The Moralistic Fallacy: On the 'Appropriateness' of Emotions." *Philosophy and Phenomenological Research* 61: 65–90.

Graham, Peter. 2010. "In Defense of Objectivism about Moral Obligation." *Ethics* 121: 88–115.

Fischer, John, ed. 1986. *Moral Responsibility*. Ithaca, NY: Cornell Press.

Fischer, John, Robert Kane, Derk Pereboom, and Manuel Vargas. 2007. *Four Views on Free Will*, edited by Derk Pereboom and Manuel Vargas. Walden, MA: Blackwell Publishing.

Foot, Philippa. 1978. "Morality as a System of Hypothetical Imperatives." In *Virtues and Vices*, 157–173. Berkeley: University of California Press.

Herman, Barbara. 1993a. "What Happens to the Consequences." In *The Practice of Moral Judgment*, 94–112. Cambridge, MA: Harvard University Press.

Herman, Barbara. 1993b. "Integrity and Impartiality." In *The Practice of Moral Judgment*, 23–44. Cambridge, MA: Harvard University Press.

Hieronymi, Pamela. 2004. "The Force and Fairness of Blame." *Philosophical Perspectives* 18: 115–148.

Hieronymi, Pamela. 2009a. "Two Kinds of Agency." In *Mental Actions*, edited by Lucy O'Brien and Matthew Soteriou, 138–162. Oxford: Oxford University Press.

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press

Hieronymi, Pamela. 2009b. "Believing at Will." *Canadian Journal of Philosophy* supplementary volume 35: 149–187.
Nagel, Thomas. 1979. "Moral Luck." In *Mortal Questions*, 24–38. New York: Cambridge University Press.
Prichard, H. A. 1949. "Duty and Ignorance of Fact." In *Moral Obligation: Essays and Lectures*, edited by W. D. Ross. Oxford: Clarendon Press.
Rosen, Gideon. 2002. "Culpability and Ignorance." *Proceedings of the Aristotelian Society* 103: 61–84.
Ross, W. D. 1939. *Foundations of Ethics*. Oxford: Clarendon Press.
Scanlon, T. M. 2008. *Moral Dimensions: Permissibility, Meaning, and Blame*. Cambridge, MA: Belknap Press of Harvard University Press.
Sher, George. 2006. "Out of Control." *Ethics* 116: 285–301.
Sher, George. 2009. *Who Knew?: Responsibility Without Awareness*. Oxford; NY: Oxford University Press.
Simmons, John. 1979. *Moral Principles and Political Obligations*. Princeton, NJ: Princeton University Press.
Singer, Peter. 1972. "Famine, Affluence, and Morality." *Philosophy and Public Affairs* 1: 229–243.
Smith, Angela. 2005. "Responsibility for Attitudes: Activity and Passivity in Mental Life." *Ethics* 115: 236–271.
Smith, Michael. 2004. "Rational Capacities." In *Ethics and the A Priori*, 114–135. Cambridge: Cambridge University Press.
Stocker, Michael. 1976. "The Schizophrenia of Modern Ethical Theories." *Journal of Philosophy* 68: 543–466.
Strawson, P. F. 1962. "Freedom and Resentment." *Proceedings of the British Academy* 48: 1–25.
Tannenbaum, Julie. 2007. "Emotional Expressions of Moral Value." *Philosophical Studies* 132: 43–57.
Tarfon. 1962. *Ethics of the Fathers*, translated by Hyman Goldin. NY: Hebrew Publishing Company.
Thomson, Judith Jarvis. 1989. "Morality and Bad Luck." *Metaphilosophy* 20: 203–221.
Thomson, Judith Jarvis. 1990. *The Realm of Rights*. Cambridge, MA: Harvard University Press.
Thomson, Judith Jarvis. 1991. "Self-Defense." *Philosophy and Public Affairs* 20: 283–310.
Watson, Gary, ed. 2003. *Free Will*. Oxford; NY: Oxford University Press.
Williams, Bernard. 1981. "Moral Luck." In *Moral Luck*, 20–39. Cambridge University Press.
Wolf, Susan. 2001. "The Moral of Moral Luck." *Philosophic Exchange* 31: 4–19.

https://doi.org/10.1080/00455091.2014.997334 Published online by Cambridge University Press