

Moral Responsibility without Wrongdoing or Blame

JULIE TANNENBAUM

Introduction

Most discussions of moral responsibility are either about what it takes to be a moral agent (i.e., the criteria for being someone to whom moral requirements apply, who can be blamed and praised, etc.) or assume the person is a moral agent and equate moral responsibility for actions such as harming or failing to aid with an agent's being blameworthy for having done wrong. This paper is situated in the latter debate, and focuses only on an individual's moral responsibility.¹ The first goal is to show that an individual can be morally responsible for harming or failing to aid even if the individual's action was not wrong, not blameworthy, not the result of blameworthy deliberation, and not the result of bad motives.² One compelling example will suffice, but since no one example speaks equally strongly to all readers, I present various sorts of examples. Moreover, these examples not only suggest how common such cases are, but also support a second goal of the paper: to establish that there is more than one way in which an individual can be morally responsible for harming and failing to aid.³

¹ I leave aside the responsibility of other kinds of agents (e.g., collective entities, individuals within a collective, etc.) and other kinds of actions (e.g., enabling).

² Of course, one can be responsible for an action without being blameworthy or having done anything wrong, such as when one brings about good outcomes for the right reasons. Sometimes, such actions are even praiseworthy. This paper focuses on responsibility for harming and failing to aid, rather than responsibility for good outcomes, and specifically, in sections 2 and 3, contexts in which the agent would have chosen a different course of action had she known or remembered before what she knows afterward.

³ Though I will drop the term "moral," I mean only to discuss moral responsibility and not any other kind of responsibility (e.g., legal). In addition, the reader can presume that what I say about harming applies equally to cases of failing to help and vice versa.

It matters whether or not we are responsible for harming, since this makes a difference to how we should relate to ourselves, and each other, in the aftermath of our action (or failure to act). For instance, it makes a difference to whether repair is called for and what sort of emotional reactions are justified. To say that an emotional reaction is justified is different from claiming it is pragmatic (useful), typical, understandable, or required.⁴ In what follows I rely on this key idea: an emotion is justified only if the belief that is conceptually associated with the emotion is reasonably believed.

I will also be assuming that when one is justified in feeling guilty, the reasonable belief conceptually associated with it is that one did the wrong thing, and not merely that there exists an unfortunate, tragic, or bad state of affairs. There are second and third person versions of this emotion, i.e., resentment by the person wronged and indignation by those who are uninvolved.⁵ These emotions are all forms of blame; to feel guilt is to blame oneself and to feel resentment and indignation is to blame others.

Some people assume that if blame is not justified (i.e., if they are not blameworthy), then they are not responsible for harming, and if they are responsible, then are blameworthy (i.e., blame is justified). And yet, if I am correct that one can be responsible for harming and failing to aid while being neither blameworthy nor having done anything wrong, then some people have blown off their responsibility when they shouldn't have, and others having been relating to themselves—or second and third parties have been relating to them—as if they were wrongdoers and blameworthy when neither is the case.

One might think, “What else is there to responsibility other than wrongdoing or blame? To be responsible for what one did to another person just is to have done the wrong thing and to be to blame (i.e., blameworthy).”

⁴ It might be typical for students to become upset at receiving a low grade, understandable for new parents to become extremely fearful when someone else holds their infant for the first time, and useful to be irate when speaking to a customer service manager. However, for a feeling (e.g., fear) to be justified, the belief conceptually associated with it (e.g., the object is a threat to something valued) must be reasonably believed. Sufficient conditions for a feeling's being justified will also take into account the feeling's intensity, duration, etc. (see Aristotle, *Nicomachean Ethics*, Book II). Lastly, if a feeling is justified, it does not follow that it is required. Being nervous before an exam may be justified but is not required; one can be fully justified in only feeling confidence.

⁵ See P. F. Strawson, “Freedom and Resentment,” *Proceedings of the British Academy* 48 (1962): 1–25.

That thought is exactly what this paper aims to challenge. Instead, I want to suggest that we should think of moral responsibility in terms of its connection to moral inadequacy: if one's action is morally inadequate (not merely not optimal), then one is morally responsible for that action. What I aim to show is that an action's being wrong is just one way that an action can be morally inadequate. The key mark of inadequacy is the need for repair: the agent in particular has something to make up for and someone to whom they need to make it up.⁶ That is, the agent has a reason, in addition to reasons that a mere bystander might have, to respond to the one who was harmed, *and* to respond in a way that counts as repair. How one makes up for what one has done is a project for another time, though the way in which one is responsible will be one of many factors in determining the kind of repair called for.

I'll begin (section I) with familiar examples of actions that are wrong and blameworthy. This sets the stage for comparisons in the remaining sections where I bring forward different types of actions, such as mere moral failures (section II) and actions with bad consequential luck (section III), that are neither wrong nor blameworthy and yet where the agent is responsible for harming or failing to aid. Nothing hangs on the terminology of "wrongdoing" or "mere moral failure," and so on; it is more for ease of reference. We could just as well refer to the actions as "type 1," "type 2," and so on. What these three types of actions have in common is that they are morally inadequate; the relationship between the agents and those harmed is now marred and thus in need of repair. It is in this sense that the agents are responsible. By distinguishing these types of actions for which we are responsible, we are able to see both the broader scope of responsibility as well as some of the different types of responsibility that have thus far been overlooked or even misconstrued.⁷

⁶ On some views of death, if the harmed person has died, repair is impossible. But the requirement for repair might nevertheless hold, without violating a reasonable interpretation of "ought implies can." See Julie Tannenbaum, "The 'Should' of Full Practical Reason," *Philosophical Books* 48 (2007): 124–35, for examples of moral requirements that are impossible to fulfill but which nevertheless do not violate a reasonable interpretation of "ought implies can."

⁷ I am not claiming that these three types of actions are the only ways in which an agent can be responsible for harming; rather they are thus far unacknowledged but important kinds of moral inadequacy.

I. Type (1) wrong and blameworthy

The first type of action (or omission) for which we are morally responsible is familiar: the agent's action is wrong and blameworthy. There are various kinds of actions that fall into this first category, such as malicious actions (done out of ill will toward another); reckless actions (the agent knowingly, but not intentionally, puts others at risk when it is unreasonable to do); and negligent actions.⁸ When agents harm due to maliciousness, recklessness, or negligence, the agent is morally responsible for their harmful action.

I will discuss morally negligent actions in detail since they will be an important contrast case for later. Here are some recognizable examples.

A student is texting a joke to a friend while driving in a highly populated area and doesn't see a pedestrian in the crosswalk. She hits and injures the pedestrian. A spouse doesn't find anniversaries and birthdays important, and so even though the spouse's partner and children make it clear that these occasions matter to them, the spouse makes no effort to remember. As a result, the spouse forgets yet again. A tenant hears a scream outside but doesn't bother to look so as not to miss the televised sports game. It turns out that a woman was being stabbed to death and would have lived had the tenant called 911.⁹

What makes these actions negligent is not merely that there is something the agents did not know, which, when combined with their action, had a bad outcome. It also matters that the agents should have known, as is illustrated by this rather contrived case.

As a homeowner enters the house, she flips on the light switch the way she usually does. But on this night, flipping the switch creates an electrical charge that flows next door and electrocutes her neighbor to death.¹⁰

⁸ This is not a complete list of type 1 actions. For example, sometimes knowingly ignoring another's will, even if one has not put him at risk (as happens in certain cases of paternalism toward someone who is autonomous), is wrong.

⁹ While some did call 911 in response to hearing Kitty Genovese's screams while being stabbed to death, others did not. Nicholas Lemann, "A call for help: What the Kitty Genovese story really means," *The New Yorker*, March 10, 2014 <<http://www.newyorker.com/magazine/2014/03/10/a-call-for-help>>. In my example, suppose there were no others to hear the screams and call 911.

¹⁰ The example is drawn from Judith Jarvis Thomson, *The Realm of Rights* (Harvard University, 1990), p. 229.

The homeowner has not been negligent (or committed any other wrong or blameworthy action). There is no reason she should have known, or even suspected, that her switch was faulty. Assume that her house was recently electrically upgraded and has been subsequently inspected. Even if this bizarre electrical problem were something she could have discovered (suppose she could have pried open her electrical system and discovered an anomaly), it is not reasonable to require the agent to go through such steps every time before flipping the switch. There is far more reason not to take apart the electrical switch than reason to do so. And so, she has not failed to take due care.

In the case of the student, however, it is reasonable to require the driver to know that texting while driving is sufficiently likely to result in harm to others on, in, or near the road. The reasons not to text (e.g., so that one is on the lookout for pedestrians) do outweigh reasons to do otherwise (barring special circumstances). The student did not take due care while driving in a populated area.

However, yet more is needed to capture the notion of moral negligence; not knowing what one should and, as a result, causing harm are insufficient for moral negligence.¹¹ It is necessary to consider *why* the agent didn't know what she should have known. Although the explanation for the student, spouse, and the tenant's ignorance is different, at bottom each lapse in knowledge reflects or results from an inappropriate orientation or disregard, either temporary or of a more permanent nature, toward others. These agents care more for themselves or their projects relative to others than they should. These three conditions together are sufficient for these agents' actions being negligent.

This point about these agents' inappropriate orientation towards others might suggest that when one is responsible for type 1 actions, what one is blameworthy for is not really the action or its results, but only the quality of one's will. Strawson, for example, might be interpreted this way when he describes blame as a reaction to someone who has not met an appropriate or justified demand for good will.¹²

¹¹ I am only interested in moral negligence and not, for example, in the definition of legal negligence.

¹² Strawson, "Freedom and Resentment," pp.76, 78, and 84. He does not specify whether one is to blame merely for the quality of one's will or also the action flowing from such a will and certain effects of that action. See also Thomas Nagel's "Moral Luck," *Mortal Questions* (Cambridge University Press, 1979), pp. 24–38, for a discussion of this issue.

However, my position departs from this in two ways. First, when I say that these agents are responsible for harming, this entails being responsible not merely for the poor quality of the will that leads to the action, but also for the action (e.g., harming) and certain of its results (e.g., the harm). So already in type 1 actions there is responsibility that extends beyond mere responsibility for the quality of one's will. The student is responsible for the pedestrian's injury, the spouse for family members' hurt feelings, and so on. The belief that is conceptually associated with the reactive attitude of blame is not (merely) about the quality of the agent's will but also, as I noted in the introduction, about the agent's having done the wrong thing.

Second, failing to meet an appropriate demand for good will is not a necessary condition for moral negligence and hence for being a type 1 action (wrong and blameworthy). Some type 1 actions result from some other blameworthy condition, such as overconfidence (unreasonable confidence).

A nurse adopts the moral ends that he should—his will is good—but he is overconfident about his memory abilities and so believes he can realize his moral ends (to aid and not harm the patient) without writing down the doctor's instructions. He later misremembers and so gives a patient the incorrect dose of medication. The patient has a seizure as a result.

The nurse is blameworthy for harming the patient because his action flows from a blameworthy condition of his, though not a problematic will. Some cases of driving while texting or tired belong here too: a driver might be overconfident in his ability to multitask while driving or in his ability to stay awake while driving in spite of his exhaustion. The trouble in these cases is not an inappropriate orientation towards others but rather an inappropriate (blameworthy, not merely mistaken) conception of oneself.

What all type 1 actions have in common is that the agents have a moral obligation, i.e., they morally should or are morally required to do or avoid some action *and* they fail to do so through a fault (blameworthy condition) of their own (e.g., overconfidence, an inappropriate orientation in their will, etc.).¹³ These agents' actions violate their obligations

¹³ Keep in mind that type 1 actions needn't have any connection to bad character. A person can be overconfident just one time.

and thus the actions are wrong. Their actions are unjustified and thus morally inadequate in this specific way. Moreover, the reactive attitude that is justified is blame (guilt, resentment, or indignation).

Type 1 action:

- The agent has a moral obligation to do x (where x might be to aid or not harm).
- The agent does not do x due to a blameworthy condition of the agent.
- The agent violates the obligation to do x.
- The agent is justified in feeling guilty.

Type 1 actions are wrong and blame is justified. But, couldn't wrongness and justifiable blame come apart? First consider whether an action could be not wrong (i.e., either permissible or required) and yet blame is justified. Perhaps, for example, blaming an innocent person maximizes the good. However, this would be a pragmatic consideration in favor of blaming, which, as I noted at the outset, is not the notion of justification with which I am working. There is also the suberogatory to consider, typically described as blameworthy action that is not wrong.¹⁴ Without taking a stand here on whether there is such a category, and whether it is yet another type of action for which we are morally responsible, we can set this category of action aside since this paper's main goal is to show that there are actions where blame is *not* justified and yet the act is wrong.¹⁵

Now consider whether an action could be wrong but for which the agent is justifiably (in a non-pragmatic sense) blameless. Below I consider type 2 and type 3 actions, which might be thought to fit this description. However, I will show that type 1–3 actions are morally inadequate in very different ways; lumping them together under the term “wrong” obscures these important differences (and leads to other problems I will discuss).¹⁶

¹⁴ See, for example, Julia Driver, “Suberogatory,” *Australasian Journal of Philosophy* 70 (1992): 286–95.

¹⁵ For the same reason, I set aside whether there are actions that can be “objectively right” but done for bad reasons or motives and hence blameworthy in this respect without being wrong.

¹⁶ Why lump recklessness and negligence together under type 1 actions? After all, the inappropriate orientation manifested by mere negligence is different from that of recklessness. Had the student believed or brought his belief to his awareness that texting was sufficiently likely to have this bad outcome he would have taken this as a sufficient reason not have texted while driving. A reckless driver, on the other hand, is aware of this but does not take it as a sufficient reason to not text and drive (though he should). While there are important differences

Of course, one can use labels anyway one likes (so long as one is clear about what one means). But in the end the issue is not semantic, but moral and metaphysical: actions can be morally inadequate in very different ways and thus we can be morally responsible in very different ways.

II. Type (2) mere moral failure

Consider these two real-life examples (though I've changed the identifying details):

Mary promises to help out her neighbor by picking her up at the airport at 6PM. In order to be sure not to be late, she heads out the door in plenty of time. 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 neighbor (who doesn't have a cell phone) has given up waiting and taken an expensive taxi ride home.

A friend is holding my child as he walks into an unfamiliar 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.¹⁷

What these agents must do is neither burdensome nor dangerous. Assume there is no more important task they should be attending to. As I see it, Mary has a moral obligation (i.e., she morally should or is morally required) to help and my friend has a moral obligation not to harm. Though these agents did not satisfy their obligations (i.e., do what they are required to do), these actions are importantly different from type 1 actions since these agents do not violate their obligation (i.e., their action is not wrong). And yet, because the agents did not satisfy their

among type 1 actions, they are all morally inadequate in this way: the agents should do (or avoid) some action and fail to do so through a fault of their own. Type 2 and type 3 actions, as I'll discuss, are not morally inadequate in this way.

¹⁷ 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 focused on one feature necessary for doing so, but who forgets to consider a second feature necessary for avoiding harm.

obligations, their actions are morally inadequate, and thus they are responsible for failing to help and harming. I'll call type 2 actions "mere moral failures." An agent merely fails just in case the agent's moral obligation is neither satisfied nor violated (and the obligation is no longer satisfiable, was not transferred or completed by someone else, etc.).¹⁸

There are various reasons one might resist my interpretation of these examples. I'll begin by first addressing whether these agents have a moral requirement to help (and not harm). Since I would like to remain theory neutral where possible, I will not pick a particular moral theory (e.g., Kantianism) and argue that these agents do have an obligation to help. Instead, I will assume that these agents have the relevant obligations unless shown otherwise.

Someone might believe otherwise due to this line of reasoning. Due their lapse in memory (e.g., not recalling the freeway construction), these agents could not come to the relevant conclusion (e.g., to take the side streets from the get go); if they could not come to the conclusion to do what would have in fact been helpful, then, we cannot saddle them with an obligation to help. Thus, they did not have any such obligation and so they are not responsible for any unmet obligation.

Implicit in this argument is the slogan "ought implies can." These agents could not help and thus they do not have an obligation to do so. While these agents were physically capable of doing what would have resulted in the recipients being helped and not harmed (Mary could have taken the side streets from the start and there was a dry spot that was not in the pathway of the lamp where my friend could have stepped), these agents were not rationally capable of doing what would have helped and not harmed.

I am assuming that sometimes a person who does not remember could have remembered. And so, these agents not only "physically can" but also "rationally can" help. Mary has no habitual way of getting to the airport or phobias that play into her choice of route. She has remembered construction alerts in the past. This time 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.

¹⁸ This is argued for in greater details in Julie Tannenbaum, "Mere Moral Failure," *Canadian Journal of Philosophy* 45 (2015): 58–84.

The same sort of story goes for my friend carrying my child. A lapse is simple just in case there is nothing of moral relevance to say about what led to the lapse. They occur, but not because one was absorbed in something else, a forgetful person who took no precautions, or prevented from being an agent at all due to stroke, and so on.

It is true that there was some state that Mary's brain was in. Does it follow from this that she could not have done otherwise? I am setting aside general concerns about free will and accepting (i) sometimes agents meet the relevant criteria of possibility (can) in the slogan "ought implies can" and (ii) sometimes agents do have obligations—obligations to act and hence also obligations to reach the decision to do that act. Merely pointing out that an agent's brain was in some state or other neither shows that the lapse was not simple nor rebuts a claim of obligation if one accepts (i) and (ii).

At this point, one might be inclined to agree that Mary and my friend were required to help and avoid harm, but think that these actions are no different from type 1 actions because these agents did not meet their obligations due to a fault of their own (and hence the actions are wrong). Mary, one might think, is like the overconfident nurse who should not have relied on his memory, or like the spouse who makes no effort to remember anniversaries and birthdays of his family members. All three should have made a backup plan (e.g., make a note of the relevant information for consultation later).

However, Mary isn't unreasonably confident. She's never forgotten a traffic alert in the past and her promise to her neighbor is no more important than other promises she has made in the past, all of which she remembered successfully without a backup plan. And Mary, unlike the spouse, does care, which is why she leaves the house on time. In addition, when the lapses are rare and the situation is not critical, backup plans—the note on the fridge, the alarm on electronic device, and so on—are not required. In fact, making a backup plan for everything one is supposed to remember is not cost free. And it can even backfire—nothing will stand out as especially important. Moreover, backup plans too are subject to mere failure. I write a note on the fridge but then forget to look at the fridge (or take the shopping list on the fridge with me to the store). It would be absurd to make a backup plan for the backup plan.

In light of this important difference between Mary's lapse and the nurse's and spouse's lapse, we can now see that there is another spouse

(again drawn from real life) who it would be a mistake to treat like the first spouse.

Spouse₂: A spouse's partner and children make it clear that birthdays and anniversaries matter to them, and the spouse, being a caring person, writes these down in a calendar with a reminder function. The reminder goes off, and the spouse looks forward to wishing the partner happy birthday, but by the time the spouse sees the partner, the spouse has forgotten.

It is not that spouse₂ doesn't care or has some subconscious desire to frustrate the partner.¹⁹ And there is no perfect time for when to set the reminder function. While spouse₂ simply forgot, that doesn't mean the spouse₂ is off the hook. Spouse₂ should have remembered and is responsible for hurting the feelings of the partner; but the spouse is not responsible in the same way as the first spouse.

The following examples are perhaps more controversial. Keep in mind that only one successful example of a type 2 action is needed, so focus on the one(s) you find most compelling.

One parent stays at work late, forgetting that the other already said that she too would be home late that night. No one is there when the children return from school. The children sit outside in the cold, hunger increasing, until one of the parents gets home.

Mark Warschauer, a professor at University of California, Irvine, sometimes took his child, Mikey, to daycare, but more often his wife did so. The route to his work and to the child's daycare overlap for much of the way. He and his wife had a tiring series of nights leading up to the morning when he offered to take his child to daycare. On the way, he forgot that he was going to daycare and forgot that his child was in the car (the child was quietly sleeping). He went to work and didn't realize that he had left his sleeping child in the backseat until it was too late. The child died from the heat.²⁰

¹⁹ Consider the refrain, "If you cared enough you would have remembered." Does this mean that we will never have a lapse when it comes to what we most care about? Such a claim requires empirical defense. Does the refrain allow that it is appropriate to care about some things more than others, and if so, does it allow that simple lapses could happen with respect to cares not at the top of our list? Or is the refrain saying that if we care just as much as we should, then we are guaranteed not to have lapses? Again, such a claim needs empirical defense. Lastly, if the refrain means that one should have cared more, such charge is not appropriate in the case of spouse₂.

²⁰ "Michael Warschauer," *KidsandCars.org*, accessed December 15, 2015 <<http://www.kidsandcars.org/michael-warschauer.html>>. His action has elicited split reactions from audience

On his website Warschauer says, "I also know there are simple techniques that might help other families avoid such tragedy. Please, please develop in your family a basic safety technique to remind yourselves that a baby is in the car seat. Put a diaper bag in the front seat every time your baby is in the car. Or put your purse, briefcase, wallet, or cell phone in the back seat where you will have to retrieve it before leaving the car." And yet on the same website his wife says, "I know this left him very tired in the daytime and may have contributed to his lack of attention. When he got that tired, he also stopped bringing his briefcase back and forth from work, thus eliminating another safety check, since when taking the briefcase out of the trunk he could see Mikey in the car seat." In this case, remembering was very important; his child's life depended on it. And he had a backup plan, as he should have given what was at stake. But his backup plan was foiled.

In Warschauer's case, he forgot because he was very tired. But suppose he was not so exhausted that he could not have remembered. There were other times when he had been this tired, had to keep track of whether he was going to work or daycare, forgotten his briefcase, but not forgotten that Mikey was in the car.²¹

What about lapses due to panic? It matters whether a person has an anxiety disorder or not, as well as whether the agent should have known this and not chosen a profession that is likely to be triggering. It also matters whether the agent has had special training to be desensitized to, avoid, or overcome panic.²² Consider this situation:

A military convoy was traveling on a road in Iraq and the second to last vehicle was blown up. All the soldiers in the last vehicle got out to help, but then came under enemy fire. One of these soldiers panicked and froze, unable to fire his weapon or reenter the vehicle. He was thrown back into the vehicle by one of his comrades and they all left the area.²³

members. Some find it hard to believe that one can forget about one's child in the car. And yet, this tragedy is not as rare as one would wish.

²¹ If he had been offered a thousand dollars that morning for getting Mikey to daycare safely, would he have remembered? Probably. Does this show he values money more than Mikey? It might only show that when things are unusual, it is easier to remember.

²² In some cases, we do require people to overcome their fear even if they have not had training.

²³ I want to thank Dr. Barbara Wettstein for sharing the military examples used in this paper. They are based on her clinical work with veterans at the West Los Angeles VA Medical Center. Identifying details have been changed.

This soldier reported feeling bad for failing to help repel the enemy fire and for endangering his fellow soldier who had to stop firing his own weapon in order to get the frozen soldier back into the vehicle. I think that the soldier is correct in thinking that he had an obligation to join in the fight and to not endanger his comrades, and that his panic doesn't rebut the claim that he has such an obligation. There was no evidence of anxiety disorder prior to or after this event and he had received special training to deal exactly with these kinds of situations. So, it is reasonable to claim that he shouldn't have panicked and should have fired his weapon (leaving aside concerns about the justice of the war).²⁴

At this point, or perhaps from the start of this section, some might agree that type 2 actions are different from type 1 actions, but, contrary to my interpretation, hold that (a) the agents are not at all responsible for harming and not helping. Others, especially after hearing about Warschauer and the soldier, might instead return to the thought I began addressing above, namely that (b) type 2 actions are not all that different from type 1 actions. I'll address each of these objections in turn.

With respect to (a), the motivating thought might be 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. On such a view, obligations to help (or not harm) are not to be taken at face value, that is, as requiring that one actually succeed in helping, but rather should be interpreted as only an obligation to *try* to help.²⁵ In the examples above the agents did try and so they satisfy the only obligations they have and thus are not responsible.

In reply, first consider how odd it is to claim that these agents only have an obligation to try given that they are, as I already discussed, physically and rationally able to do more. More importantly, to treat these agents' actions as success stories fails to morally distinguish them from real moral success stories, e.g., Sue, who promised to pick up her neighbor at the airport and remembered the freeway construction, and so took the side streets from the start. She arrived on time and brought

²⁴ In this paper, the soldier's action is the only mere failure that is not a result of forgetting or not recalling a fact necessary for taking the correct means. I discuss a variety of other kinds of mere failures in "Mere Moral Failure."

²⁵ Other interpretations include: an obligation to adopt the end of helping or an obligation to act with the intention of helping.

her neighbor safely home. The neighbor's getting picked up on time was not simply a bonus.

To make even clearer how different the actions are of someone who has good intentions and fails, as opposed to someone who has good intentions and succeeds, notice that Mary, not Sue, has a reason to apologize, and not merely explain why she was so late. If a common friend were to later explain what happened (how Mary had set out on time but forgotten about the freeway construction), Mary wouldn't be off the hook; she still needs to follow up with the neighbor herself, even if all the parties know that this information was communicated. If all Mary were obligated to do is try, which she did, then there would be nothing for her to apologize for. And yet there is.

Turning to (b), are these agents' actions type 1 actions?; did they violate their obligation to help or not harm and would blame be justified? After all, as I said, they do have something to make up for and someone to whom they need to make it up. And yet, these agents care, and in no way exhibit an inappropriate orientation towards others, or any other blameworthy condition that characterizes negligent actions, or any other type 1 actions. They are wholly aimed at helping and not harming those they should; the soldier, for example, jumps out of the vehicle ready to provide covering fire for others. They do not have, for example, optimistic planning bias; they are not lazy, and so on.

This might lead some to think instead that type 2 actions are wrong and blameworthy—just less so than type 1 actions. Blame should be mitigated when it comes to type 2 actions.

The concept of mitigation deserves a closer look than I can provide in this paper, but here are some initial comments.²⁶ Mitigation seems to be driven by facts that tend to undercut the agency of the agent, such as when one is overcome by strong emotions, cognitively overloaded, and so on. But these kinds of considerations also undermine the claim that there was an obligation (in light of "ought implies can"). I argued above, the rational and physical agency of type 2 actions is not undercut—this is

²⁶ A difference in degree of blame relative to the harm done is not, to my mind, a kind of mitigation (e.g., compare a texting driver that injures to a texting driver that kills). How much harm was done is relevant to determining what level of blame is justified in the first place. In any case, the difference between type 1 and type 2 actions is not the degree of harm, as spouse₁ and spouse₂ make clear.

most plausible in the case of Mary, my friend, and spouse₂. They neither are overcome by strong emotions nor cognitively overloaded. The agency of Warschauer and the soldier was under pressure. And I am assuming that what they must do is hard, given the psychic strain they are under. But not impossible. And its being hard doesn't itself explain why blame would be mitigated. It is often hard to resist doing what is wrong but that alone doesn't mitigate blame for wrongdoing (e.g., when cheating on an exam etc.), so why would it do so in these cases? And why would that only partially rather than wholly mitigate blame?

So, are type 2 actions examples of (completely) blameless wrongdoing? I agree the agents are not worthy of any blame (blame, to any degree, is not justified), but I resist the suggestion that they've done wrong since these agents do not fail to help due to a fault of their own. If we insist on labeling type 2 actions as wrong, then conceptual connections among a triad of moral concepts (wrong, blameworthy, violation of obligation) is lost.²⁷ Also, this seems to be only a semantic triumph (one can now claim that there is no responsibility without wrongdoing). It misses the spirit of the paper, which is to point out that there are many types of responsibility for harming other than type 1 actions, which tend to be the main focus of philosophers and non-philosophers alike. Pointing out that type 1 and type 2 actions can be brought under a common label obscures the important differences between them.

More importantly, categorizing type 2 actions as blameless wrongdoing leads to a contradiction. If type 2 actions were wrong (but blameless), then the agents would be justified in having some sort of critical emotional response to their actions or omissions. What would the thought content associated with their feeling be? If the thought associated with their emotional reaction were, "I did the wrong thing," then their feeling would be no different from guilt (self-blame). And so, type 2

²⁷ One might keep the triad by claiming that type 2 actions are blameworthy but no one is justified in blaming the agents because we all lack the standing to do so (see Susan Wolf, "The Moral of Moral Luck," *Philosophic Exchange* 31 [2001]: 4–19). Their failures are all too familiar. We ourselves have messed up in similar enough ways and it would be hypocritical of us to blame them. But notice that it is not only inappropriate for others to blame them, but also inappropriate for type 2 agents to blame themselves. Guilt is as much out of order as is resentment or indignation, given that the agents do not fail to help due to a fault of their own. Thus, both the blameworthiness of these agents' actions and blaming in all its forms (guilt, resentment, and indignation) is not justified.

agents would be justified in blaming themselves.²⁸ Thus, those interested in labeling type 2 actions as blameless wrongdoing must face up to the challenge of explaining how the feelings these agents are justified in having, which involves placing themselves on the hook of moral responsibility, differs from that of self-blame.²⁹

There do not seem to be any other categories into which one can easily slot type 2 actions.³⁰ So, we are left with understanding type 2

²⁸ Some take guilt to involve not only the thought that one has done a wrong act and a feeling (feeling bad about what one has done), but also the thought that one deserves to feel bad (see, for example, Derk Pereboom “Responsibility, Regret, and Protest,” *Oxford Studies in Agency and Responsibility*, vol 4, ed. David Shoemaker [Oxford University Press, 2017]). And so, one might try to distinguish the feelings of type 1 and type 2 agents by claiming that type 2 agents do not have this extra thought, while type 1 agents do. This is a controversial interpretation of guilt, but in any case, an explanation is owed as to why it would be inappropriate for type 2 agents to have the thought that they deserve to feel bad for what they have done.

²⁹ A similar challenge might apply to those who would describe these agents’ action as objectively wrong but subjectively right (Peter Graham, “In Defense of Objectivism about Moral Obligation,” *Ethics* 121 [2010]: 88–115), but it would depend on the criteria of objective wrongness and subjective rightness (recall that the right thing to do was rationally accessible to type 2 agents).

The challenge facing a view like Scanlon’s is different. Blame, according to Scanlon, is not tied per se to wrongdoing (T. M. Scanlon, *Moral Dimensions: Permissibility, Meaning, and Blame*, [Belknap Press of Harvard University Press, 2008]). An act is wrong if and only if disallowed by principles that no one could reasonably reject; it does not matter what the agents’ reasons for acting were (pp. 23 and 124). Blame, on the other hand, 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, it would seem, contrary to common sense, that the nurse in section 1 (type 1 action) cannot justifiably blame himself since he does not have 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 type 2 actions do not have any special meaning for the agents’ relationships with those harmed. And yet the moral relationship between agents who merely fail and those they failed is impaired (repair is called for). So, either type 2 actions are wrong, but (contrary to Scanlon) wrongdoing itself does have moral meaning, or (contrary to Scanlon) the meaning of an action is not exhausted by the agent’s attitude to others.

³⁰ Type 2 actions are not examples of the category Sarah Buss calls “justified wrongdoing” (“Justified Wrongdoing,” *Nous* 31 [1997]: 337–69). Buss’s examples and type 2 actions have this in common: the agents are blameless and yet accountable. But her agent is blameless because justified, whereas type 2 actions are not justified. Type 2 agents can be reasonably expected to know both what would help and not harm whereas in Buss’s examples we “suspect [the agent] was justifiably ignorant” (338). Moreover, they are not what Elizabeth Harman calls “mere moral mistakes” (“Morally Permissible Moral Mistakes,” *Ethics* 126 [2016]: 366–93), which are failures to do what is supererogatory (not morally required) and that are criticizable. Lastly, invoking the category of pro tanto wrongdoing won’t help us make sense of type 2 actions. This language is typically used for justified actions with a moral remainder (e.g., you cross my lawn without my permission to save someone’s life; what you do is all things

actions as neither wrong nor blameworthy, but instead as mere moral failures. If blame is not justified, what emotional reaction is? We don't have a name for this feeling, as is the case for many feelings. In spite of lacking a natural language term, it is a type of feeling bad for what one did; one feels one failed the other person(s). The thought conceptually associated with this feeling is that one didn't do what one should have; the thought is not that the action was wrong.

Without an awareness of the difference between type 1 and type 2 actions, there is a tendency of those who merely fail to both mischaracterize their actions as wrong and their feelings as guilt.³¹ Friends, family, and therapists often correctly resist an agent's blaming herself, but, because they also tightly associate blame and responsibility, they also tend to incorrectly judge the agent as not responsible for failing to help. Many therapists view any feelings and expressions of responsibility as distortions that the patient needs to be freed from.³² Sadly, therapy for soldiers like the one discussed above is too often unsuccessful. Perhaps some patients are less likely to adopt the therapist's view of their actions and feelings—this is extremely speculative—because the therapist's views do not speak to, or help them make sense of, what the soldiers know on some level is true, namely, that they are responsible.

Here is a quick summary of the features of a type 2 action.

Type 2 action:

- The agent has a moral obligation to do x (where x might be to aid or not harm).
- The agent does not do x through *no fault* (blameworthy condition) of the agent.

considered justified, but some think that you ought to apologize because what you did was pro tanto wrong). But type 2 actions are not justified all things considered.

³¹ In "Emotions in 'Black and White' or Shades of Gray? How We Think about Emotion Shapes Our Perception and Neural Representation of Emotion" (*Psychological Science* 27 [2016]: 1428–42), Ajay B. Satpute et al. demonstrate that when given two extreme options for describing an emotion (e.g., as calm vs. fearful), rather than a more scalar or nuanced set of choices, individuals shifted to whichever extreme they were hovering closest to. In the case of self-assessment, they became more prone to have the corresponding feeling. One can see how this might extend to individuals given only a choice of thinking of responsibility in terms of blaming or not.

³² These empirical claims are based on limited interviews with social workers, psychologists, and psychiatrists.

- The agent neither violates nor satisfies the obligation to do x.
- The agent is justified in feeling that she failed the other person(s).

Type 2 actions are morally inadequate in this way: the agent merely failed to do what she was morally required to do. The agents are morally responsible for harming and failing to aid. But because the agents have all the regard and respect for others that they should, and no blameworthy condition of theirs led to their action, their action is neither wrong nor blameworthy. Mere moral failures are not morally inadequate in all the ways the type 1 actions are, and thus, they are not unjustified; but neither are these agents' actions justified.³³

III. Type (3) bad moral luck

There is yet a third category of actions where one is responsible for one's action, though in a quite different way from the previous two types described. Again, I'll begin with an example.

Darin Strauss was driving the speed limit with due care down a four-lane highway where people sometimes walked and rode bikes on the shoulder of the highway. Two seniors from his school were riding bikes on the shoulder. He was aware of them and then suddenly one of them turned her bicycle into his lane; he hit and killed her.³⁴

In 2005, I was told a somewhat similar story by Maryann Gray, who earlier in her life had killed a child who ran in front of her car.³⁵ Their actions bring to mind Bernard Williams's example of the lorry driver, who through no fault of his own runs over a child.³⁶ Both Darin and

³³ Bernard Williams draws a distinction between an action's being justified, not justified, unjustified in "Moral Luck," *Moral Luck* (Cambridge University Press, 1981), pp. 20–39, though my usage is somewhat different from his.

³⁴ See "Life after Death," *This American Life*, episode 359, July 18, 2008 and Darin Strauss's memoir, *Half a Life* (Random House, 2011).

³⁵ For a recent interview of Maryann, see Alice Gregory, "The Sorrow and the Shame of the Accidental Killer," *The New Yorker*, September 18, 2017, pp. 28–32. Maryann eventually established a website to help those who have accidentally killed others: "Coping with Causing a Serious Accident: A Site for Information, Support, and Healing" <<http://www.accidental-impacts.org/>>.

³⁶ See "Moral Luck," p. 28.

Maryann suffered mentally for decades after these killings. Psychologists claim that drivers who have done nothing wrong and yet harmed others are at greater risk of developing post-traumatic stress disorder than those who have done something wrong and as a result harmed another.³⁷ Darin's friends and parents tell him that he's not blameworthy, and the parents of the student he killed originally took this position too, but then their emotions took a negative turn and they sued him. The students at his school took sides for and against Darin. I suspect that people who blame Darin (including Darin himself at times) do so because they don't know what else to reach for in order to make sense of the thought or feeling that he is more than causally responsible for the student's death.³⁸ And those who are confident that blame is not justified rush to the conclusion that he is not morally responsible for harming her.

Williams claims that these drivers are justified in feeling what he labels "agent-regret." This feeling is different from mere regret, which anyone who was not involved would be justified in feeling. Mere regret is not a feeling of responsibility but rather the feeling one has in response to thinking about things of value (people, animals, art) that are harmed or destroyed—even when this happens as a result of natural events like tidal waves. Conceptually associated with mere regret is the thought that it would have been better had this not happened, and the wish that things had gone differently, even if they could not have gone differently. It need not involve seeing anyone as morally responsible in any way. It is sadness in response to the loss of value.

However, agent-regret is a feeling of responsibility. What agents regret is their own agency. They see the harm they've caused as something that can be laid only at their feet, and that nothing that an insurance program or others could do would alter or address the fact that the agents in particular need to repair the relationship.³⁹ The belief

³⁷ See D. L. Delahunty et al., "Acute and Chronic Distress and Posttraumatic Stress Disorder as a Function of Responsibility for Serious Motor Vehicle Accidents," *Journal of Consulting and Clinical Psychology* 65 (1997): 560–7 and E. Hickling et al., "Effects of Attribution of Responsibility for Motor Vehicle Accidents on Severity of PTSD Symptoms, Ways of Coping, and Recovery over Six Months," *Journal of Traumatic Stress* 12 (1999): 345–53.

³⁸ See note 31.

³⁹ This sets it apart from legal notion of strict liability, which is understood as a form of legal responsibility that may be covered by insurance. Strict liability is legal responsibility for

that is conceptually associated with agent-regret is the moral inadequacy of one's action. If the feeling of agent-regret is to be justified, the thought that Darin's action is morally inadequate must also be justified. But what is morally inadequate about his action?

Unlike type 1 actions, Darin did not kill due to a fault of his own. He has no inappropriate orientation towards others or blameworthy self-conception that led to his action. He adopted the ends he should have (and not any he shouldn't have), his beliefs were reasonable even if incorrect (he reasonably but falsely believed he could drive down that street at that time with that speed without killing anyone), he drew the rationally supported conclusion from his deliberations, and he acted on this conclusion. He had mistaken information and thus was ignorant, but his ignorance is not blameworthy and thus his action is not negligent. And his action is certainly not reckless or malicious or any other type 1 action. His action is not morally inadequate in a type 1 way.

And unlike type 2 actions, Darin did not fail to satisfy any obligations that he had. After the fact, Darin realizes (or at least has reason to believe) that there was an alternative course of action, which, had he taken it, would not result in his killing the student. Had Darin driven down a different street or highway, or driven more slowly, or driven more quickly, etc. then he would not have killed the student. But while these alternative courses of action were physically available to him, they were not rationally accessible to him through no fault of his own. And so, Darin had no obligation to avoid killing the student in that way. That is, it is not the case that he should have known better or should have acted otherwise.

What, then, is morally inadequate about Darin's action?⁴⁰ To see our way to an answer, first consider how an action is evaluated when morality is not at issue. Suppose I value great tennis. There are many ways my value might manifest itself: I might adopt the end of watching professional

damages, or injury, "even if the person found strictly liable was not at fault or negligent" ("Strict Liability," *West's Encyclopedia of American Law*, edition 2, 2008, retrieved May 8, 2017 <<http://legal-dictionary.thefreedictionary.com/strict+liability>>). Also, the motivation for strict liability seems to be to incentivize agents (people, institutions, companies, etc.) to take extra care or to distribute the costs of injury to the party best situated to absorb them, whereas these pragmatic considerations are not part of the moral notion of responsibility with which I am working.

⁴⁰ The following material is drawn from Julie Tannenbaum, "Emotional Expressions of Moral Value," *Philosophical Studies* 132 (2007): 43–57.

tennis, or not interfering with people who are playing well, or playing tennis well myself. There is no requirement that my value manifest itself in one of these ways rather than another. However, suppose I do adopt the end of playing tennis well.

Here are some ways in which my tennis playing can be justifiably judged as unsuccessful. Suppose that on a windy day I'm inattentive to how the wind will affect my shot, and so the ball goes out. In this case, my action is unsuccessful in part because the decision that led to it is faulty. If playing tennis well is my end, then I should, rationally speaking, adopt the relevant sub-ends, like being attentive to what will help or hinder the ball landing in the court. My action is also unsuccessful if I'm a bit lazy and so fail to act on my decision to hit with more force. Now a more complicated example: suppose I'm lazy and so fail to hit the ball with the force it is reasonable to think it needs to get over the net, but a sudden and unpredictable wind carries the ball over the net and into the court. My action is nevertheless inadequate. It is not an instance of playing tennis well, but merely a lucky shot.

What if my decision is flawless in the sense that it flows from my adopting and pursuing the end and sub-ends related to playing tennis well and I act on my decision (e.g., exhibit no weakness of will)? In some situations, in spite of my best intentions, effort, and attention, I hit the ball out due to a slight increase in the wind that could not have been reasonably expected (given the weather forecast, the location of the court, etc.). This action is also not a success. There is no do-over; my opponent is awarded the point, not me. My body does what I instructed it to do, but given the wind, the results are not what I intended (the ball goes out). I cannot find fault with my shot insofar as I look at my will (practical reasoning). But insofar as the shot fails to realize a *realizable end* that I set for myself (playing tennis well), my action is inadequate. Merely going for an end and not realizing it is not enough to make an action inadequate. It matters that there was an alternative shot I could have made that would have landed in even given the wind (a cross court shot with a lot of topspin). That is why agent-regret involves, as Williams says, the thought of how much better it would have been had I acted otherwise *and* the thought that I could have acted otherwise. Even though this shot was not rationally accessible, it was physically accessible.

These tennis examples indicate that there are at least these two necessary conditions for successful action. First, my action must be

done on the basis of good practical reasoning (having a consistent set of ends, adopting the rationally entailed sub-ends, willing those sub-ends or actions that are instrumental to, or constitutive of the sub-ends, etc.). Second, my action must realize my realizable ends and sub-ends. Perhaps I have more control in meeting the first condition, but nevertheless, both are equally relevant to the evaluation of my action as successful. For an action to fully embody one's ends or what one values, the action must flow from, that is, be the intentional realization of, the agent's realizable ends.

Whether one values tennis, and how that value shapes one's ends, is optional. However, when it comes to morality, there are certain ends one is morally required to adopt (the malicious agent fails here) and weigh in one's deliberation in specific ways (the reckless agent fails here) and one is required to adopt certain sub-ends of due care in light of those ends (the negligent agent fails here). Type 3 agents succeed in all these ways, and thus type 3 actions are not unjustified. In addition, type 3 agents form all the relevant beliefs they are rationally required to form (unlike type 2 agents) and meet all the moral obligations that it is reasonable to place upon them (unlike in type 2 agents). But, as in the tennis example, type 3 agents fail to realize their moral ends/sub-ends, which are in fact realizable, and thus their actions are not fully justified. One of Darin's moral (sub)ends is to drive without harming anyone. Darin chose a course of action that did not realize this end (he killed the student⁴¹) when an alternative was available that would have realized this end. His action is inadequate as an action, since he fails to embody his end, and it is *morally* inadequate, since he fails to embody his moral end.

Type 3 actions are morally inadequate neither because of a mere mismatch between the intended end and the result, nor because they have bad results that matter morally. After all, in the trolley case where one switches the trolley headed towards five onto a track with one, and in

⁴¹ The way Darin is causally connected to the death of the student makes his action a killing. This is not true of his friend who was a passenger. What about the student; does she kill herself? Maybe, but that doesn't show that Darin doesn't also kill her. Moreover, even if she is morally responsible for killing herself, that does not preclude Darin from also being morally responsible for killing her. Also, keep in mind that doing that without which the death would not have resulted is not sufficient to make one a killer. This is a point about the metaphysics of actions, not morality. For example, if a doctor initiates and then removes life support that is not self-sustaining, he does not kill the patient, but lets the patient die, even though, but for the doctor's withdrawal of aid, the patient would not have died (Jeff McMahan, "Killing, Letting Die, and Withdrawing Aid," *Ethics* 103 [1993]: 250-79).

medical triage cases where some are aided and not others, there are bad results that matter morally. In the trolley and triage cases the agents' actions are justified (assuming the agents have moral ends, reasonable and accurate beliefs about their unfortunate options, etc.). Their actions do embody their moral values. The agents, insofar as they are virtuous, will regret that circumstances led to such a difficult and tragic choice, but the agents do not regret the action they choose; they do not wish that they had chosen an alternative course of action. But type 3 agents do wish they could have chosen differently, precisely because their actions do not embody their moral values. Type 3 agents fail to realize their moral ends with the course of action they chose.

Here are two further examples of type 3 actions inspired by real-life events.

A doctor has good reason to think a patient requires penicillin. The patient has taken it in the past without adverse side effects. However, this time the patient develops Stevens-Johnson Syndrome and as a result of extensive tissue damage on the body and eyes, the patient is permanently scarred and blinded in one eye.

John sees a fellow soldier acting in an unsettled way, which is not uncommon after a mission. John is not particularly close to the soldier but they are in the same unit. John thinks about whether to talk with the soldier and see if he's doing all right, but decides against it. He doesn't want to be too intrusive. The behavior the soldier is exhibiting is not uncommon, and usually after blowing off some steam, things are back to normal the next day. But that night the fellow soldier goes into the armament and kills himself.

Analysis similar to that of Darin's action applies to the doctor's action, which demonstrates the wide range of contexts in which type 3 actions occur (consider also interactions between police and citizens). Recognition of type 3 actions can open up a new space of mutual and self-understanding without mistakenly heading into the territory of blame.

John's situation is more ambiguous. If he should have talked with the soldier, then my analysis would imply that his action was a mere failure (type 2 action) rather than a bad luck case (type 3 action). But it is not clear that he should have done so, given that it can be intrusive or disrespectful (by sending the message that the other is seen as weak or incapable). Also, talking with every soldier who seems unsettled after a mission would be overwhelming for John, especially while holding

himself together after a mission. Seen in this light, John didn't fail with respect to any of his obligations to his fellow soldier. And yet, John expressed that he wished he had talked with the soldier and also believes that had he talked with him, he would have averted the suicide.⁴² John was morally required to have the end of caring for his comrades, and he did not realize this end though he could have. If that is true, then his action is type 3—not a type 3 harming, but a type 3 failure to aid.

In summary, here are the key features of type 3 actions:

- The agent does *not* have an obligation to do x (where x might be to aid or not harm) because any such obligation is not rationally accessible to the agent.
- The agent satisfies any obligations that the agent has.
- The agent does not realize a realizable end that the agent is morally required to adopt.
- Feelings of agent-regret are justified.

The condition of realizing one's realizable ends is not a moral requirement but a metaphysical (necessary, not sufficient) condition for fully successful action. But not realizing a realizable *moral* end makes an action *morally* unsuccessful. This type of moral inadequacy is what underpins the claim that Darin, the doctor, and John are responsible for harming and failing to aid. Perhaps agents like these, and others who kill and fail to aid without any wrongdoing, suffer for so long because it is hard to make sense of the feeling of responsibility alongside the thought that they are not blameworthy. On the account I have offered, they are correct in thinking they are responsible. Their mistake was in assuming that there is only one type of responsibility, namely, the blameworthy kind.⁴³

⁴² Sometimes, the agent, bystanders, and those who later hear about the case lack sufficient evidence for such a belief.

⁴³ Hanna Pickard also distinguishes responsibility from blame (see "Responsibility without Blame: Empathy and the Effective Treatment of Personality Disorder," *Philosophy, Psychiatry, and Psychology* 18 [2011]: 209–23, and her other closely related papers). According to Pickard, (i) to be responsible just is to exercise agency—that is, for an agent to know what she is doing, choose to do it, and act on her choice; and (ii) to be morally responsible for harm, for example, is to be blameworthy, that is, to be responsible and lack an excuse (pp. 212–15). Detached blame just is a judgment of blameworthiness (210), while affective blame is feeling entitled to negative emotions because of what the blamed person has done (218). Her main point is that (i)

Conclusion

If we accept that there is moral responsibility not only in type 1, but also in type 2 and type 3 cases, it will be important to delineate the scope of these types of actions and to consider the ways in which this makes a difference to how one goes about repair. We should also not assume that all cases will exhibit only one type of responsibility.

Even if the above is eventually accomplished, it will nevertheless often be hard to know what sort of action we, or others, have done. That is, it may be difficult to determine whether there really is an obligation to help or not harm and, if so, whether one has violated it, merely failed with respect to it, or met the obligations one had but failed instead to realize one's realizable values in one's chosen course of action (or omission). And hindsight will not always clarify the matter.⁴⁴

What we can be clear about is that we should widen our moral focus beyond that of blame and wrongdoing. We need not reach for blame as our only way of registering that an agent is morally responsible for what they've done.⁴⁵

does not entail (ii), nor does it conceptually or psychologically entail affective blame. But this framework won't help us when it comes to the cases in this paper. In type 2 and 3 cases, and even negligent actions (a type 1 case), the agents do not meet (i), since they do not know that they are harming until after the fact, and so they neither meet (ii) nor merit affective blame. And yet, in at least some cases of negligent action both forms of blame are in fact justified and all negligent agents are blameworthy and morally responsible. Moreover, I reject (ii) and thus distinguish moral responsibility from being blameworthy, but without appealing to an account of responsibility like (i). As I see it, if one's action is morally inadequate, then one is morally responsible for that action. All three types of actions are morally inadequate, though the actions exhibit different kinds of moral inadequacy, and hence different kinds of moral responsibility; two of the three kinds of moral responsibility do not involve being blameworthy.

⁴⁴ One might claim that the type 2 and 3 agents are justified feeling bad just in case it is reasonable for them to be uncertain about whether their action is a type 1 action, otherwise they have nothing to feel bad about. I have tried to show that there is something for such agents to feel bad about, namely the moral inadequacy of their action. But even if one rejects this, there is no reason to accept the idea that uncertainty makes one justified in feeling bad. If I'm not sure whether, for example, I have benefited or harmed you, I don't see that I have any reason to feel one way or the other.

⁴⁵ I want to thank Barbara Wettstein, Nishi Shah, and members of the audience at the Workshop in Normative Ethics at the University of Arizona, Practical Philosophy Workshop at the University of Chicago, Bled Philosophical Conference in Slovenia, Center for Human Values at Princeton University, and Grand Rounds at the West Los Angeles VA Medical Center.