

*Moral Injury and Moral Failure**Lisa Tessman****Introduction**

Military service – particularly combat – puts service members at a high risk of having an experience of moral failure of a kind that few civilians face: the experience of having to perform an action – often killing someone or letting someone die – that one realizes is the best thing that one can do in the situation, but that nevertheless strikes one as being morally wrong or even unthinkable. Grasping the unthinkability of an action that one has to perform is an emotional experience of extreme aversion and distress. The experience may result in what is called moral injury, namely, in the words of one veteran, “the kinds of wounds we inflict on ourselves that come inherently with the wounds we inflict upon others.”¹ Killing another person – especially if it is someone who is particularly vulnerable, such as a child – or not saving a person whom one loves or for whom one is responsible are among the primary actions that cause moral injury.

Moral injury is considered to be *moral* because of the real or perceived moral wrongdoing that is at its source and because it is characterized by certain moral attitudes – attitudes of holding oneself responsible or taking

* I first presented an earlier draft of this chapter at the Ethics of War and Peace Conference, which was held at the United States Military Academy, West Point, in October 2019; I also presented versions of it at the Institute for Advanced Studies in the Humanities at Binghamton University in February 2020; at the Central New York Moral Psychology Workshop in November 2020; at the Jackson Family Center for Ethics and Values, Coastal Carolina University, in February 2021; and as the Sullivan Lecture at Fordham University, in April 2021. I benefitted enormously from comments from participants at these events. I am grateful to graduate students in my Spring 2020 seminar on Moral Responsibility, who read and discussed a draft. I would also like to thank individuals who commented on earlier drafts, including: Mich Ciurria, Barrett Emerick, Laura Engel, Ami Harbin, Alex King, Katie Stockdale, and Mark Wilson. Most of all, I thank my spouse, Bat-Ami Bar On, for contributing to my thoughts in this chapter in many ways; as a veteran herself, she knew a lot about living in the aftermath of war. She died unexpectedly on November 16, 2020.

¹ Tyler Boudreau, “The Morally Injured,” *The Massachusetts Review*, 52, no. 3/4 (2011), 748.

oneself to have failed morally, and feelings of guilt about this failure. And moral injury is called an *injury* because it often manifests in debilitating symptoms and self-destructive behaviors that, at the more extreme end, include withdrawal, abuse of drugs and alcohol, and self-harm, including suicidal ideation and attempts at suicide.² There are mental health practitioners who specialize in treating moral injuries, and researchers who investigate what causes moral injury, how it can be prevented, and how the symptoms can be more effectively treated. This is all important work, and I believe that service members and veterans are owed much better treatment than what is currently available. But this is not my focus. Instead, I am interested in what service members and veterans who experience moral injury thereby express about moral failure. I want to salvage the “tragic truth” that I believe they often express – namely that some of the actions that it may be beyond our control to avoid are not just wrong but unthinkable – and that tends to get dismissed because of the discrepancy between their own attitudes of holding themselves responsible, and other people’s attitudes of releasing them from responsibility.³

Experiences of Moral Injury

According to the model of moral injury developed by some of the leading psychiatrists and psychologists in the field, the experiences that may cause moral injury include:

perpetrating, failing to prevent, bearing witness to, or learning about acts that transgress deeply held moral beliefs and expectations... Moral injury requires an act of transgression that severely and abruptly contradicts an individual’s personal or shared expectation about the rules or the code of conduct, either during the event or at some point afterwards.⁴

² Sonya B. Norman and Shira Maguen, “Moral Injury,” National Center for PTSD, www.ptsd.va.gov/professional/treat/cooccurring/moral_injury.asp.

³ As Peter Marin writes about this “tragic truth”: “Our great therapeutic dream in America is that the past is escapable, that suffering can be avoided, that happiness is always possible, and that insight inevitably leads to joy. But life’s lessons... teach us something else again, something that is both true of, and applicable to, the experience of the vets. Try as they do to escape it, the past pursues them; the closer they come to the truth of their acts, the more troubled they are, the more apart they find themselves, and the more tragic becomes their view of life.” Peter Marin, “Living in Moral Pain,” in Robert Emmet Meagher and Douglas Pryer (eds.), *War and Moral Injury: A Reader* (Eugene, OR: Cascade Books, 2018 [1981]), 178.

⁴ Brett Litz, et al., “Moral Injury and Moral Repair in War Veterans: A Preliminary Model and Intervention Strategy,” *Clinical Psychology Review*, 29, no. 8 (2009), 700. Psychiatrists and psychologists who treat moral injury clinically and who research the causes and treatments of moral injury include Brett Litz, Shira Maguen, and William Nash. Like them, I am focused on cases in which the injury comes from one’s own moral failure. However, Litz, Maguen, Nash, and others’ conception

The U.S. Department of Veteran Affairs recognizes that “In the context of war, moral injuries may stem from direct participation in acts of combat, such as killing or harming others, or indirect acts, such as witnessing death or dying, failing to prevent immoral acts of others, or giving or receiving orders that are perceived as gross moral violations.”⁵ A series of studies led by Shira Maguen has revealed that killing in war is highly associated

of moral injury is a departure from the understanding of moral injury that was introduced by psychiatrist Jonathan Shay; see Jonathan Shay, *Achilles in Vietnam: Combat Trauma and the Undoing of Character* (New York: Simon and Schuster, 1994); and Jonathan Shay, *Odysseus in America: Combat Trauma and the Trials of Homecoming* (New York: Scribner, 2002). According to Shay, the injurious situations include those in which the injured party is injured not by perpetrating or being complicit or associated with wrongdoing, but rather by being betrayed or wronged by someone in authority. Using Achilles’ betrayal by Agamemnon in the *Iliad* as his model, Shay describes moral injury as a “betrayal of what’s right” by someone in authority (Shay, *Achilles in Vietnam*, p. 3). Shay himself distinguishes his definition of moral injury from that of Litz, Maguen and Nash, noting that the term “has been used in two related, but distinct, senses; differing mainly in the ‘who’ of moral agency. Moral injury is present when there has been (a) a betrayal of ‘what’s right’; (b) either by a person in legitimate authority (my definition), or by one’s self – ‘I did it’” (Litz, Maguen, Nash, et al.); (c) in a high stakes situation”; Jonathan Shay, “Moral Injury,” *Psychoanalytic Psychology*, 31, no. 2 (2014), 182. Shay adds that “there’s not a lot we can do to eliminate the sort of Moral Injury that Nash, Maguen, and Litz have written about”; “Moral Leadership Prevents Moral Injury,” in Robert Emmet Meagher and Douglas Pryer (eds.), *War and Moral Injury: A Reader* (Eugene, OR: Cascade Books, 2018 [2010]), 304. Shay believes that given his own definition, moral injury can be prevented through better military leadership, because it is betrayal on the part of the leadership that causes moral injury. He also points out that betrayal by the leadership – for instance, by giving commands that require immoral actions – is what often leads to service members’ own moral failures, and thus to moral injury as defined by Litz, Maguen, Nash and others (*ibid.*, 305). Loosely following Shay, some scholars employ a very broad understanding of moral injury; Nancy Sherman, for instance, writes that moral injury “refers to experiences of serious inner conflict arising from what one takes to be grievous moral transgressions that can overwhelm one’s sense of goodness and humanity. The sense of transgression can arise from (real or apparent) transgressive commissions and omissions perpetrated by oneself or others, or from bearing witness to the intense human suffering and detritus that is a part of the grotesquerie of war and its aftermath. In some cases, the moral injury has less to do with specific (real or apparent) transgressive acts than with a generalized sense of falling short of moral and normative standards befitting good persons and good soldiers”; Nancy Sherman, *Afterwar: Healing the Moral Wounds of Our Soldiers* (New York: Oxford University Press, 2015), p. 8. See also Nancy Sherman, *The Untold War: Inside the Hearts, Minds, and Souls of Our Soldiers* (New York: W.W. Norton and Company, 2010). I am examining a narrow rather than broad range of situations that can cause moral injury. While I believe that moral injuries that result from one’s own moral failure and those that result from the failures of leadership are both important, there are crucial differences in the emotions or moral attitudes that are associated with each, and they should not be lumped together. This is confirmed by Schorr et al. in a study of what veterans perceive as having been morally injurious in their military experience: “Based on our qualitative results, we posit that an important organizational framework for considering morally injurious events is the veteran’s attribution of responsibility for the event. An event in which he or she feels responsible for what occurred (either through an act of omission or commission) is likely to result in different emotional responses, and thus have different implications for treatment than events in which the responsibility is attributed to others (e.g. guilt and shame versus anger and frustration); Yonit Schorr et al., “Sources of Moral Injury Among War Veterans: A Qualitative Evaluation,” *Journal of Clinical Psychology*, 74, no. 12 (2018), 12.

⁵ Norman and Maguen, “Moral Injury.”

with a number of different mental health and functional consequences, which together can be taken to characterize moral injury. Furthermore, the severity of the effects was found to depend on the type of killing that the veteran had engaged in; to put a rough gloss on it, those who have killed civilians exhibit more severe symptoms than those who have killed only enemy combatants, and those who have killed women, children, and/or the elderly, as well as those who have killed prisoners, suffer more than those who did not.⁶ It is not simply killing, or killing of a certain kind, but also the veteran's attitude toward their own killing that matters; the attitude of "combat guilt" is strongly associated with, for instance, suicidal ideation and suicide attempts.⁷ Retrospective thoughts or memories of their own acts of killing – particularly but not only forms of killing that were not justified – are often accompanied by guilt, shame, a feeling of having failed morally, or a judgment that one is a bad person. As Noah Pierce, an Iraq War veteran, wrote in his suicide note, "I am not a good person, I have done bad things. I have taken lives, now it's time to take mine."⁸ The concept of moral injury is meant to capture a linked cluster of experiences (such as killing), moral emotions (such as guilt and shame), and harmful effects; this cluster differs from the cluster that characterizes the more well-known phenomenon of post-traumatic stress disorder (PTSD), which typically includes life-threatening experiences, the primary emotion of fear, and symptoms such as hypervigilance.

⁶ Shira Maguen et al., "The Impact of Killing in War on Mental Health Symptoms and Related Functioning," *Journal of Traumatic Stress*, 22, no. 5 (2009), 435–43. In a study of Vietnam combat veterans, those who had killed had significantly higher rates of "PTSD symptoms, dissociation experiences, functional impairment, and violent behaviors," even after controlling for other kinds of combat exposure (*ibid.*, 440). A later study of Vietnam veterans found that those with "killing experiences were twice as likely to report suicidal ideation as those who did not kill, even after accounting for PTSD, depression, substance use disorder diagnoses, and adjusted combat exposure"; Shira Maguen et al., "Killing in Combat May Be Independently Associated with Suicidal Ideation," *Depression and Anxiety*, 29 (2012), 921. In a similar study of Iraq War veterans, Maguen and her colleagues again found a strong relationship between killing – whether directly or indirectly – and "PTSD symptoms, alcohol abuse, anger, and relationship problems"; Shira Maguen et al., "The Impact of Reported Direct and Indirect Killing on Mental Health Symptoms in Iraq War Veterans," *Journal of Traumatic Stress*, 23 (2010), 86. A study of Gulf War veterans found that "taking a life in combat is a critical ingredient in the development of postdeployment mental health concerns"; Shira Maguen et al., "The Impact of Killing on Mental Health Symptoms in Gulf War Veterans," *Psychological Trauma*, 3, no. 1 (2011), 24. And a study of both Iraq and Afghanistan veterans found that those "with particular killing experiences involving non-combatants and/or anger or revenge are most at risk"; Shira Maguen et al., "Killing and Latent Classes of PTSD Symptoms in Iraq and Afghanistan Veterans," *Journal of Affective Disorders*, 145, no. 3 (2013), 348.

⁷ Herbert Hendin and A.L. Haas, "Suicide and Guilt as Manifestations of PTSD in Vietnam Combat Veterans," *The American Journal of Psychiatry*, 148, no. 5 (1991), 586–91.

⁸ Ashley Gilbertson, "The Life and Lonely Death of Noah Pierce," *The Utne Reader*, February 13, 2009. www.utne.com/politics/lonely-death-noah-pierce-ptsd-iraq-war.

Killing is not the only act that is experienced, in the context of military operations, as a moral failure. Other moral transgressions may include acts such as torturing or humiliating prisoners or detainees, destroying property, sexually assaulting civilians or fellow service members, injuring others either intentionally or accidentally, using disproportionate or unnecessary violence, acting cruelly, or ordering others to commit a variety of actions, as well as omissions such as failing to protect a fellow service member or a civilian, being negligent in a way that results in harm or death, or not intervening when witnessing someone else's immoral behavior. While all of these acts can result in a moral injury to the service member who commits the act, some cases differ from others in an important respect: in some cases, the act is fully avoidable, while in other cases, there is no way that the service member could have avoided doing what they did without doing something worse or equally bad. Many cases lie somewhere in between the avoidable and the unavoidable. I will be focusing primarily on unavoidable failures, with the understanding that it is frequently hard to know whether or how a particular act could have been avoided; often what haunts a service member or veteran is the thought that perhaps they could have done something better than they did. When an act cannot be avoided, at least not without doing something worse or equally bad, then the moral status of the act is complicated by the fact that one can say both that it was the best possible thing to do and *in that sense* was the right thing to do, while in some other sense it was still a wrongdoing.⁹ But before getting to unavoidable moral wrongdoings, I will say a few words about those that are avoidable.

Consider cases in which a service member tortures or humiliates prisoners for entertainment, rapes a civilian or another service member, or uses more force or violence than is necessary for accomplishing a task. For example, think about the following story that a military chaplain told to journalist David Wood, about a soldier who:

...had been in a firefight and had his rifle sight trained on an insurgent, but instead of firing at the man's chest – the “center of mass” that troops are taught to aim for – he consciously lowered his weapon and shot him in

⁹ Most moral philosophers would deny that there can be such a thing as an impossible-to-avoid moral wrongdoing, because they accept the principle that “ought implies can” – that is, that one cannot be morally required to do anything that one cannot do. I have argued that “ought implies can” does not apply to all moral requirements and that there indeed are impossible moral requirements and unavoidable moral failures. See Lisa Tessman, *Moral Failure: On the Impossible Demands of Morality* (New York: Oxford University Press, 2015); and Lisa Tessman, *When Doing the Right Thing Is Impossible* (New York: Oxford University Press, 2017).

the stomach, knowing that would cause a lingering and painful death. As an army medic worked to save the man, the soldier and his squad members gathered around and watched the man die in agony.¹⁰

The soldier came some months later to the chaplain because he was suffering from the anguished feelings characteristic of moral injury; the soldier had realized: “I’m not just killing because I have to, but because I want this to be a suffering,” and he knew that to be wrong. A painful violation of his own moral values.¹¹

In this case, the soldier could have and should have avoided the wrongdoing. The distress he experienced afterward is entirely appropriate. In fact, it is a sign that he has at least some moral conscience. When veterans self-reflectively blame themselves for intentional, avoidable moral wrongdoings, they are right to take responsibility for their actions. Some therapeutic approaches have assumed that moral injury (like PTSD) is a matter of being plagued by *distorted* cognitions – such as the belief that one is guilty of wrongdoing when in fact one is not – and that the therapist should aim to change the distorted self-blaming attitudes. Other mental health professionals contend that this assumption about distorted cognitions is misguided, precisely because there are cases in which the veteran has committed a moral wrongdoing and *properly* grasps this. For instance, Gray, Nash and Litz point out that “individuals can and do occasionally act intentionally in ways that violate their values and standards of conduct, and... such actions – when reflected upon rationally – may give rise to significant guilt, shame and distress... In cases in which at least partial culpability is real and rational, the assignment of blame to oneself or others must also be rational, appropriate, and accurate.”¹² They argue that in such cases it is a mistake for the therapist not to recognize that self-blame is warranted.¹³ It is fairly clear that veterans who react to their own avoidable wrongdoings by taking responsibility are accurate when they take themselves to have failed morally, though of course they might take either too much or too little responsibility for an action for which they are only

¹⁰ David Wood, *What Have We Done: The Moral Injury of Our Longest Wars* (New York: Little, Brown and Company, 2016), p. 21.

¹¹ *Ibid.*

¹² Matt Gray, William Nash, and Brett Litz, “When Self-Blame is Rational and Appropriate: The Limited Utility of Socratic Questioning in the Context of Moral Injury,” *Cognitive and Behavioral Practice*, 24 (2017), 384.

¹³ Their approach, called Adaptive Disclosure, thus assumes that: “(a) pain means hope – anguish, guilt, and shame are signs of an intact conscience and self-and-other-expectations about goodness, humanity, and justice; (b) goodness is reclaimable over the long haul; and (c) forgiveness (of self and others) and repair are possible regardless of the transgression”; *ibid.*, 386.

partially to blame. When veterans hold themselves responsible for wrongdoings that they could have and should have avoided committing, they are likely to see their own response to themselves – self-blame – reflected in a parallel reaction from other people, who may also blame the veteran.

However, this tends not to be the case when veterans hold themselves responsible for having committed some terrible act that was at the same time the best action they could have taken in the circumstance – they will not see their reaction to themselves reflected in other people's reactions, for other people will have a tendency *not* to hold them responsible. In these sorts of cases veterans may have committed the terrible act under some degree of compulsion, or in the face of a moral dilemma in which all available options would have involved wrongdoings, or in situations in which they lacked crucial information for making a decision, or lacked sufficient control to do what would have been right. Alternatively, they may have committed the act because the only alternative would have been extreme self-sacrifice. In other words, the action that the veterans understand to have been wrongdoings are not necessarily actions that they could have avoided, at least not without doing something worse or equally bad, or not without self-sacrificing more than anyone expects them to. For instance, consider this case, related by David Wood, in which a marine probably did the best thing that he could have done in the situation:

Nikki...was twenty-two years old when he was sent as a marine infantryman to Afghanistan, where he shot and killed a young boy. This was not uncommon in the murderous confusion of our recent wars, where farmers and mothers and young kids might seize a weapon and shape-shift in a moment into a combatant and back again to an innocent civilian, and young Americans peering into the murk would have a moment to decide: kill or not. This time, an exhausting firefight with Taliban insurgents had dragged on for hours across the superheated desert wastes and tree-lined irrigation canals of Helmand Province. Late that afternoon, Nik saw from the corner of his eye someone darting around the corner of an adobe wall, spraying bullets from an assault rifle held against his small hips. Nik swiveled his M4 carbine, tightened his finger on the trigger, and saw that it was a boy of maybe 12 or 13. Then he fired. According to the military's exacting legal principles and rules, it was a justifiable kill, even laudable, an action taken against an enemy combatant in defense of Nik himself and his fellow marines. But now Nik is back home in civilian life, where killing a child violates the bedrock moral ideals we all hold. His action that day, righteous in combat, nonetheless is... a painful violation of the simple understanding of right and wrong that he and all of us carry subconsciously through life.¹⁴

¹⁴ Wood, *What Have We Done*, pp. 7–8.

In Nik's own words: "He was just a kid. But I'm sorry, I'm trying not to get shot and I don't want any of my brothers getting hurt, so when you are put in that kind of situation...it's shitty that you have to, like...shoot him. You know it's wrong. But...you have no choice'.¹⁵ Nik recognized that he faced a moral dilemma: he had to choose between the wrongdoing of killing a child and the wrongdoing of not protecting his fellow marines.

It is not always a moral dilemma that makes moral failure unavoidable. Sometimes the failure occurs because of mistakes due to unavoidable ignorance – for instance, ignorance about whether an approaching individual is an innocent civilian or an insurgent. David Wood's story about two marines, Doss and Canty, illustrates this:

There was an Afghan kid who used to come around the marines' outpost, basking in their attention. One day he kept trying to get them to go outside, wanting to show them something. Doss and Canty and some others finally went with him into an adjacent field toward a tree line and discovered a desiccated corpse, a man the marines had shot a while back... The man had been acting suspiciously and wouldn't respond to the warning shots the marines had fired in accordance with the ROE [rules of engagement], and so he was shot and killed. Now his body had been ravaged by dogs. It turned out that the corpse was the boy's father, who was deaf and mute and couldn't hear or respond to warning shots, the boy explained, and that's why he appeared to ignore the marines and why he was shot and killed and was now lying dead in a field.¹⁶

One of the marines, Doss, whose own father had long been ill and who died two years after Doss' return from Afghanistan, began to have nightmares in which it was his own father whose body was lying half eaten away in the field.

Although there are particular incidents like these in which something that is easily granted to be morally horrifying – such as killing a child, or killing a vulnerable civilian – is unavoidable, the type of killing that is most clearly sanctioned in war – that is, the killing of an enemy combatant – may also be experienced as a moral wrongdoing, and killings of this type are unavoidable as long as war is unavoidable.¹⁷ Killing enemy

¹⁵ *Ibid.*, pp. 13–14.

¹⁶ *Ibid.*, p. 140.

¹⁷ As a study on moral injury by Schorr et al. notes: "Injuring and killing the enemy in battle is central to the core mission of war. Service members are trained to kill and depart for war with the expectation that this is what they are setting out to do. And yet, a number of veterans...described being changed by taking these actions that violated or, at the very least, challenged their previously held moral beliefs"; Schorr et al., "Sources of Moral Injury Among War Veterans," 6.

combatants might at the same time achieve one morally good aim – if the aim of the mission is itself morally justifiable, carrying out the mission necessitates killing, and the killing is conducted in accordance with humanitarian laws – while violating other moral values. As Yonit Schorr et al. put it:

an individual may be able to justify his actions given the context within which the event occurred (e.g. self-defense or defense of others) and may even acknowledge having honored one set of morals by this action (defense of country, following orders, protecting his men), and yet still experience considerable distress about how those actions conflict with a separate set of morals (e.g. do not kill, protect the weak).¹⁸

Thus the simple fact of moral value pluralism – namely, the existence of multiple, irreducibly different moral values – and the inevitability of conflict between the moral requirements that these different values entail is enough to make moral failure unavoidable. It is not always possible for service members to simultaneously fulfill all of the nonnegotiable moral requirements that they face.¹⁹ Because war requires killing, and because each human life that is lost is of the kind of irreplaceable value that we tend to take ourselves to be morally required not to sacrifice, war regularly makes moral failure inescapable, even for those who always perform the morally best action that is possible in the circumstance.²⁰

In the case of *avoidable* wrongdoing, other people, including psychotherapists, may readily find it appropriate for veterans to see themselves as having failed morally. What about in the case of unavoidable wrongdoing? Veterans often still feel an anguished sense of responsibility for what they take to have been their own moral failures, even when they recognize that those failures were unavoidable. However, when the failure was unavoidable, or could have been avoided only by doing something worse or equally bad, or through extreme self-sacrifice, other people do not tend to hold them responsible despite the fact that the veterans hold themselves

¹⁸ Ibid., 6.

¹⁹ For my full discussion of moral value pluralism and its relation to unavoidable moral failure, and for my distinction between negotiable and nonnegotiable moral requirements, see Tessman, *Moral Failure and When Doing the Right Thing is Impossible*.

²⁰ For an important discussion of the relationship between moral value pluralism and veterans' expressions of moral failure, see Tine Molendijk, "Toward an Interdisciplinary Conceptualization of Moral Injury: From Unequivocal Guilt and Anger to Moral Conflict and Disorientation," *New Ideas in Psychology*, 51 (2018), 1–8. See also Tine Molendijk, Eric- Hans Kramer, and Désirée Verweij, "Moral Aspects of 'Moral injury': Analyzing Conceptualizations on the Role of Morality in Military Trauma," *Journal of Military Ethics*, 17 (2018), 36–53.

responsible.²¹ This is an interesting mismatch of reactions. It might seem that the mismatch should be resolved by assuring veterans that they are not responsible for their unavoidable actions,²² but I argue that this denies something important that their distress reveals.

I believe that both veterans' tendencies to hold themselves responsible and other people's reluctance to hold them responsible get something right.²³ The veterans may be right that they are responsible in the sense of having failed to meet a binding moral requirement, even when the requirement was impossible to meet. This makes their tendency to hold themselves responsible a *fitting* response because it correctly represents their own action as a wrongdoing. However, others may at the same time be right that the veterans are not responsible in the sense of being blameworthy for having violated shared normative expectations that are fair. To blame others for certain failures, including those that are unavoidable, is inappropriate for a specific reason: it reflects normative expectations that cannot be shared expectations, because one cannot reasonably expect others to accept that certain actions – including those that are impossible to perform – are required. But we may (rightly) have expectations of ourselves that are not shared expectations and that do not have to be reasonable. My claim depends on acknowledging both the *fittingness* of taking oneself to be responsible for an unavoidable failure, and the *unfairness* of the normative expectation that one avoid the unavoidable. I will suggest

²¹ The mental health professionals (such as Gray, Nash, and Litz) who deny that the anguished emotions characteristic of moral injury are necessarily “distorted cognitions,” and who advocate accepting rather than challenging the veteran’s sense of responsibility for the action, seem to leave open the possibility that actions that are unavoidable in the context of war may still be appropriately conceived of by the veteran as moral failures. However, Gray, Nash, and Litz do argue in support of their approach by emphasizing that “moral injury can and often does result from instances of intentional perpetration” (Nash and Litz, “When Self-Blame is Rational and Appropriate,” 383), so it seems that they have *avoidable* actions in mind, and do not commit one way or the other to the appropriateness of taking responsibility for *unavoidable* actions. See also Litz et al., “Moral Injury,” 695–706; and William Nash and Brett Litz, “Moral Injury: A Mechanism for War-Related Psychological Trauma in Military Family Members,” *Clinical Child and Family Psychology Review*, 16, no. 4 (2013), 365–75.

²² Nancy Sherman’s chapter in this volume aims to find resources in the reactions of others to help those suffering from moral injury. The approach I take here explains how self-guilt for unavoidable actions can be fitting and that acknowledgment by others of this self-guilt can be an important part of the recovery process.

²³ In some ways, this is similar to Susan Wolf’s position, though she is focused on cases of moral luck rather than specifically on cases of unavoidable moral wrongdoing, so she has no reason to say, for instance, that the agent is right in taking himself to have committed a wrongdoing; rather, her point is that there is a virtue that has to do with taking responsibility for one’s actions when, due to luck, they have a bad outcome. See Susan Wolf, “The Moral of Moral Luck,” *Philosophic Exchange*, 31, no. 1 (2000), 1–16. Thanks to Barrett Emerick for pointing out the connection with Wolf’s thinking.

that the phenomenon of moral injury should be addressed in a way that both acknowledges taking responsibility as a fitting response and recognizes the value of having shared normative expectations that are fair.

Reactive Attitudes

Work on what are called “reactive attitudes” offers a helpful framework for understanding moral responsibilities. Work in this tradition takes our sentiments – and specifically the sentiments that we experience in reaction to how other people regard or disregard us and how we regard or disregard others – and our associated practices of holding each other, and ourselves, responsible to be the basis for our responsibilities. I accept this overall approach but find one aspect of it to be problematic: existing work on the topic has for the most part denied that we can be responsible for unavoidable moral failure, and it has pushed to the side the reactive attitudes that are indicative of this kind of responsibility.²⁴ I think that there is a sense in which we can be responsible for unavoidable failures, and failures that would require great self-sacrifice to avoid, even when others cannot rightly hold us responsible for them, and that a case can be made for this by rethinking the implications of the reactive attitudes that are experienced as moral injury.

Reactive attitudes include both *self-reactive attitudes*, which are the attitudes a person has in response to the quality of their *own* regard for others, and *interpersonal reactive attitudes*, which are reactive attitudes that people have in response to the quality of *another* person’s regard for ourselves or for others. In typical cases in which a moral injury results from an unavoidable action, there is a discrepancy between the reaction that a

²⁴ Similarly, work on moral injury has pushed the self-reactive attitudes to the side by suggesting that veterans are wrong to hold themselves responsible for many of their combat-related actions. As Camillo Bica noted in one of the earliest discussions of moral injury, “veterans are advised to ignore what has occurred, to ‘deresponsibilize,’ i.e. to neutralize their feelings by accepting the ‘naturalness’ of their behavior on the battlefield...”; Camillo C. Bica, “A Therapeutic Application of Philosophy: The Moral Casualties of War: Understanding the Experience,” *International Journal of Applied Philosophy*, 13, no. 1 (1999), 82. An exception is Krista Thomason who discusses the guilt that former child soldiers experience even while other people (such as mental health professionals) assure them that they are not responsible for what they did. In this case, what leads people to exempt the child from moral responsibility is the very fact that by virtue of being a child, they are not considered a full moral agent. Furthermore, even if the child is considered to be a moral agent, the fact that they were coerced to become a soldier is thought to excuse them from responsibility. Nevertheless, Thomason argues that the child’s guilt is both intelligible and appropriate in the sense of being morally valuable. Krista Thomason, “Guilt and Child Soldiers,” *Ethical Theory and Moral Practice*, 19, no. 1 (2016), 115–27. Thanks to Katie Stockdale for directing me to Thomason’s paper.

moral agent (in this case, a service member or veteran) has to herself (the self-reactive attitude) and the reactions of other people to the moral agent (the interpersonal reactive attitudes). The typical self-reactive attitude in these sorts of cases is to feel distress in the form of guilt, shame, and self-blame, while the typical interpersonal reactive attitude is not any kind of blaming response; the distressed self-reactive response seems to reflect the way that it matters to the moral agent that they have morally failed (even if unavoidably), while the absence of an interpersonal blaming response seems to reflect an understanding that the moral agent had no morally better option, or that choosing the morally better option would have been supererogatory (that is, good but not required), so no shared normative expectation for them to have acted differently could be a fair expectation. There is a lack of symmetry between the typical self-reactive attitude and the interpersonal reactive attitude; I believe that *both* types of reactions can be appropriate, and thus that it is right for there to be an asymmetry. My aim is to make sense of the asymmetry by distinguishing between different ways in which a reactive attitude can be appropriate.²⁵

I will begin with a brief review of the way that Peter Strawson characterized the reactive attitudes when he first introduced the concept in 1962.²⁶ Taking interpersonal reactive attitudes (as opposed to self-reactive attitudes) as the paradigm case, he notes that we have natural reactions – in the form of sentiments, or attitudes – to how other people regard and treat

²⁵ This is a kind of self-other asymmetry that Michael Slote has argued is absent from at least a certain version of virtue ethics. See Michael Slote, "Morality and Self-Other Asymmetry," *The Journal of Philosophy*, LXXXI, no. 4 (1984), 179–92; and *From Morality to Virtue* (New York: Oxford University Press, 1992). I disagree with Slote about this because I deny that the self-other asymmetry is due to a vice related to an improper balance between self-regarding and other-regarding traits. I don't think that the asymmetry is necessarily due to any shortcoming. Rather, I argue that we may *rightly* hold ourselves responsible for things that no one else can hold us responsible for. See Lisa Tessman, "The Virtues of Reactive Attitudes," *The Journal of Value Inquiry*, 55 (2021), 437–56. Thanks to Nancy Sherman for making the connection between Slote's work on this topic and my claims in this chapter.

²⁶ P.F. Strawson, "Freedom and Resentment," *Proceedings of the British Academy*, 48 (1962), 187–211. Strawson develops the concept in the context of a discussion of determinism – a context that I will leave to the side. Strawson argues that accepting "the truth of a general thesis of determinism" (*ibid.*, 195) would have no effect on our practices of responsibility – it would not stop us from having reactive attitudes that express how and in what way other people's regard for us (and for others) matters to us. It would not lead us to excuse people from responsibility for their actions (where we tend to excuse when a bad action turns out not to have been based on any underlying ill will), and it would not lead us to exempt all people from counting as responsible agents (as, for instance, we would exempt certain categories of people, such as young children) and thus to see people as "an object of social policy...as a subject for what...might be called treatment...as something...to be managed or handled or cured or trained" (*ibid.*, 194). Instead, regardless of the truth of determinism, the

us, that is, to their attitudes toward us.²⁷ They are called *reactive* attitudes because they are attitudes in reaction to attitudes, from the perspective of those engaged as participants in interpersonal relationships. Reactive attitudes contain a mixture of affective and cognitive elements; they lie somewhere between how we feel in response to what we take someone's regard to be and how we judge someone's regard. To capture this mixture, I refer to reactive attitudes as a kind of *evaluative response*. The reactive attitudes present someone's regard for others as having certain evaluative features; for instance, my resentment might present the quality of someone's regard for me as insulting or slighting, and my gratitude might present the quality of someone's regard for me as kind or benevolent. It is because it matters to us how other people regard and treat us that, when we engage with them, we tend to react in positive or negative ways to their quality of regard for us; these reactions reflect the ways in which their regard matters to us. Strawson emphasizes "how much we actually mind, how much it matters to us, whether the actions of other people – and particularly of *some* other people – reflect attitudes towards us of goodwill, affection, or esteem on the one hand or contempt, indifference, or malevolence on the other."²⁸ Our interpersonal reactive attitudes, in reflecting the fact that it matters to us how others regard us, serve as a "certain sort of demand for inter-personal regard."²⁹ The fact that regard matters gives us a constitutive reason to have a particular reactive attitude, namely one that makes a demand for regard. This is what connects *matter*ing with *holding responsible*.³⁰ Our reactive attitudes "rest on, and reflect, an expectation of, and demand for, the manifestation of a certain degree of goodwill or regard

fact that we are committed to having interpersonal relationships in which other people's quality of will matters to us means that while it would be logically possible to adopt this "objective attitude" towards others rather than responding to them as a participant in an interpersonal relationship (i.e. responding with the reactive attitudes), it is "practically inconceivable" (ibid., 197) that we would do so; our valuing of interpersonal relationships precludes it.

²⁷ Theorists in the Strawsonian tradition have refined what Strawson calls "quality of will" into several different aspects of our wills. For instance, David Shoemaker argues that we sometimes have a reactive attitude in response to someone's "quality of character," sometimes in response to their "quality of judgment" and sometimes in response to their "quality of regard"; David Shoemaker, *Responsibility from the Margins* (New York: Oxford University Press, 2015). Given this division, my focus is on quality of regard.

²⁸ Strawson, "Freedom and Resentment," 191.

²⁹ Ibid., 201.

³⁰ This characterization borrows from Pamela Hieronymi's account, in which she distinguishes constitutive reasons from extrinsic reasons; Pamela Hieronymi, "The Wrong Kind of Reason," *The Journal of Philosophy*, CII (2005), 437–57. See also Hieronymi's notion of "responsibility as mattering," which she presents in several places, including in her 2020 unpublished manuscript "Introduction" to *Minds That Matter*, available at: <https://hieronymi.humspace.ucla.edu/in-progress/>.

on the part of other human beings towards ourselves; or at least on the expectation of, and demand for, an absence of the manifestation of active ill will or indifferent disregard.”³¹

Given our “human commitment to participation in ordinary interpersonal relationships,” and given that in such relationships it matters to us how others regard us, it would be “practically inconceivable” for us to give up our reactive attitudes and thus our practices of holding others responsible.³² In this conception, our attitudes and our practices of holding responsible are antecedent to our responsibilities; they are the (or at least a) source of our responsibilities. As Gary Watson explains: “It is not that we hold people responsible because they *are* responsible; rather, the idea (our idea) that we are responsible is to be understood by the practice, which itself is not a matter of holding some propositions to be true, but of expressing our concerns and demands about our treatment of one another.”³³ This is an important point because it means that the answer to the question of whether or not we are responsible for certain moral failures (including those that are unavoidable) is to be found in our reactive attitudes – in our practices of holding responsible that reflect what matters to us, not in some independently existing fact about what responsibilities we have or what is morally required of us.

However, not every reactive attitude establishes what we are “really” responsible for, because some of our reactive attitudes are simply inappropriate responses that we should dismiss rather than take to have any normative implications for us. We need to consider what the criteria are for a reactive attitude to be appropriate. If self-reactive attitudes of feeling responsible for moral failure are inappropriate (as we might think that they are either because failure was unavoidable, or because our only alternative would have been supererogatory), then perhaps they are not indicative of our having failed to meet any real responsibilities.

What makes a reactive attitude appropriate or inappropriate? To begin with, the reactive attitudes are appropriate only when they “fit” what they are reactions to. For a reactive attitude to fit its object, we must correctly apprehend the features of the object – whatever it is we are reacting to – and our reaction must be evaluative, such that it reflects the way in which the object of our reaction really does or does not matter to us; because a reactive attitude is an evaluative response, what it reflects when it fits

³¹ Strawson, “Freedom and Resentment,” 200.

³² *Ibid.*, 197.

³³ Gary Watson, *Agency and Answerability: Selected Essays* (Oxford: Clarendon Press, 2004), p. 222.

its object are the evaluative features of the object. Thus, I could have an unfitting reaction either by making a mistake about what the object is, or by being somehow mistaken about what really matters. The first kind of mistake is straightforward: if I take you to have ignored me and I resent you for ignoring me, then in order for my resentment to be appropriate, I must be correct that you have ignored me; if it turns out that you weren't ignoring me – you simply didn't see me – then my resentment will be inappropriate. Because the same action could be performed with quite different underlying attitudes, the response to an action will be an inappropriate – unfitting – reactive attitude whenever one makes a wrong assumption about the underlying attitude. Take Strawson's example:

If someone treads on my hand accidentally, while trying to help me, the pain may be no less acute than if he treads on it in contemptuous disregard of my existence or with a malevolent wish to injure me. But I shall generally feel in the second case a kind and degree of resentment that I shall not feel in the first.³⁴

Once someone's intention in performing an action becomes apparent, it may (though in some cases may not) serve to either justify or excuse what would otherwise have appeared to be a wrongdoing. We would think someone *justified* in treading on my hand if they did it in order to squash the spider on my hand; we would *excuse* them for treading on my hand if they did not even see that my hand was there, because they were focused on swatting the spider off of my head; either way, they had no ill will, so if ill will is what matters – which, in Strawson's view, it is – then a reactive attitude of resentment would be inappropriate.³⁵

One might thus conclude that any reactive attitude that serves to hold someone responsible for performing an *unavoidable* action – no matter how terrible the action is – is inappropriate. It may seem that whenever an agent could not have avoided a particular action, it is inappropriate to hold them responsible, simply because they have exercised no ill will. Their actions will seem to be either justified or excused. For instance, if a soldier kills an enemy combatant because doing so was necessary for meeting an overriding moral requirement to pursue the aims of a war (that was itself morally justified), their action may (assuming it meets various other

³⁴ Strawson, "Freedom and Resentment," 191.

³⁵ As J.L. Austin said of justifications and excuses, when we have a justification for our action, "we accept responsibility but deny that it was bad" and when we have an excuse for our action, "we admit that it was bad but don't accept full, or even any, responsibility"; J.L. Austin, "A Plea for Excuses," *Proceedings of the Aristotelian Society*, New Series, 57, no. 1 (1956–57), 2.

criteria, such as being in accordance with the rules of engagement) count as *justified*. If a soldier kills an innocent civilian whom they unavoidably mistook to be posing a threat, their action may be *excused*.³⁶ Thus interpersonal reactive attitudes that *forbear* from holding veterans responsible for unavoidable actions – even when the actions are terrible – seem to get it right; they seem to be fitting.

Furthermore, it is not just for entirely *unavoidable* actions that it may be appropriate to forbear from holding other people responsible by accepting justifications or excuses. When an action is unavoidable, we recognize that it was not performed out of ill will or lack of regard of any kind. But often what matters to us about other people is that they regard others with what we take to be *enough* good will – we do not necessarily expect that their concern for others will always override all other considerations, including a healthy regard for themselves. If someone chooses not to sacrifice their life in order to avoid perpetrating some terrible act, such as killing a child, others may be inclined to excuse their choice or see it as justified. When people do not blame veterans for acting in self-defense in such situations, it reflects a sense that it is unreasonable to expect others to make certain self-sacrifices. It matters to us that someone acts with concern toward others, but our normative expectations in interpersonal relationships allow the participants to also exercise some level of concern for themselves. Our typical interpersonal reactive attitudes serve to constitute certain levels of self-sacrifice as supererogatory. Just as we tend not to hold other people responsible when they fail to do the impossible, in many cases we also do not hold them responsible when they choose not to make extreme self-sacrifices. There will, however, be a limit to this. There are acts so terrible – atrocities – that sacrificing one's own life to avoid being a perpetrator may be morally required rather than supererogatory. And service members – by virtue of accepting their role in the military – are already committed to making sacrifices that in other contexts would be seen as supererogatory.

³⁶ In addition to justifications and excuses, there are also *exemptions* from responsibility. We *exempt* someone from responsibility if they are simply incapable of being responsible, for instance, because they are a young child, or an adult who is disabled from agency either permanently (such as someone who suffers from dementia) or temporarily. What counts as having one's own agency compromised for reasons beyond one's control can be difficult to determine. For instance, consider a service member who is driven by sudden and uncontrollable rage to take cruel revenge on an enemy combatant who has just killed their closest friend. Some people might tend to say the service member was in a state in which they could not be responsible for their actions; other people may be reluctant to grant this, and the service member himself may rightly resist being treated as a nonagent.

This tendency to release other people from responsibility stands in contrast to what research into moral injury has shown: that service members or veterans frequently hold themselves responsible for unavoidable actions and actions that would have entailed what other people take to be supererogatory self-sacrifice. It turns out that it is really just our *interpersonal* reactive attitudes – our reactions to *other* people – that follow a pattern of not holding people responsible for unavoidable or difficult-to-avoid failures. In the same sorts of situations, the typical *self*-reactive attitude is a feeling of responsibility for the action. However, theorists have tended to dismiss self-reactive attitudes as misguided whenever they diverge from interpersonal reactive attitudes. It seems that the guilt, shame, and self-blame that are characteristic of moral injury would, in this view, be deemed inappropriate self-reactive attitudes, and thus attitudes that do not tell us anything about what we “really” are responsible for. Is this right? Should we just dismiss the self-reactive attitude that consists of a feeling of responsibility for what are understood as unavoidable moral failures or failures to do what others take to be supererogatory?

I will suggest that we should not dismiss self-reactive attitudes just because they do not match what we agree are fitting interpersonal reactive attitudes. Rather, I will argue that there are differences between the criteria by which self-reactive attitudes should be judged to be fitting and the criteria by which interpersonal reactive attitudes are judged to be fitting, because self-reactive attitudes may fittingly be either a subjective evaluative response or an evaluative response that reflects shared normative expectations, while interpersonal reactive attitudes can only be fitting when they reflect what can reasonably be expected to be shared values or shared normative expectations. However, before examining the differences in what makes self-reactive and interpersonal reactive attitudes *fitting*, I will turn in the next section to another possible way in which one might object that the self-reactive attitudes that are typical in cases of moral injury are inappropriate, for “inappropriate” does not always mean “unfitting.”

Right and Wrong Kinds of Reasons

It is possible for a reactive attitude to be fitting but to be inappropriate in some other way; for instance, a reactive attitude may be called inappropriate in the sense that experiencing it is harmful or inexpedient. As Justin D’Arms and Daniel Jacobson have pointed out, the term “appropriate”

tends to be applied in ways that conflate fittingness with other concerns.³⁷ This “conflation problem”³⁸ is evident, I believe, in the charge that it is inappropriate to hold oneself responsible for unavoidable moral failures or failures to do what others take to be supererogatory. Distinguishing fittingness from other concerns will allow us to see that the self-reactive attitudes that are typical in cases of moral injury may be fitting despite there being reasons of other kinds – reasons based on other kinds of concerns – not to hold oneself responsible for unavoidable moral failures or failures to do what others take to be supererogatory. Reasons that bear on the question of what an object’s evaluative features are are the “right kind of reasons” for counting the object as actually having those evaluative features; other kinds of reasons that weigh in favor of or against having an evaluative response are the “wrong kind of reasons” because they are merely *extrinsic* reasons – that is, they do not bear directly on the question of whether or not an object has those evaluative features – instead, they may be moral or prudential reasons for or against having a particular evaluative response.³⁹ An attitude fits its object when the reasons for calling it

³⁷ Justin D’Arms and Daniel Jacobson, “The Moralistic Fallacy: On the ‘Appropriateness’ of Emotions,” *Philosophy and Phenomenological Research*, 61, no. 1 (2000), 65–90. They argue that “an emotion can be fitting despite being wrong to feel” for either moral or prudential reasons; *ibid.*, 65. As they put it elsewhere, “to call a response ‘appropriate’ is vague praise... only certain good reasons for or against having a response bear on the associated evaluative judgment”; Justin D’Arms and Daniel Jacobson, “Sentiment and Value,” *Ethics*, 110, no. 4 (2000), 731. What is important about this for them is that it poses a problem for theories such as (neo)sentimentalism that rely on what they call the “response dependency thesis,” namely the thesis that “to think that X has some evaluative property Φ is to think it appropriate to feel F in response to X”; *ibid.*, 729. The problem is that it may be inappropriate to feel F in response to X for reasons that have nothing to do with whether or not X has the evaluative property Φ , so it is impossible to know whether or not X has the evaluative property Φ simply on the basis of whether or not it is “appropriate” to feel F.

³⁸ *Ibid.*, 732. This problem also came to be called the “wrong kind of reasons problem.” See Wlodek Rabinowicz and Toni Rønnow-Rasmussen, “The Strike of the Demon: On Fitting Pro-Attitudes and Value,” *Ethics*, 114, no. 3 (2004), 391–423.

³⁹ It is Pamela Hieronymi who introduces the terminology of “constitutive reasons” (to designate what had been called “right kinds of reason”) and “extrinsic reasons” (to designate what had been called “wrong kinds of reasons”), and it is she who discusses the “wrong kind of reasons problem” in terms of how a reason bears on a question. She argues that we should reject the definition of a reason as “a consideration that counts in favor of an action or attitude” (because this definition does not allow us to distinguish right from wrong kinds of reasons) and instead “think of a reason as a consideration that bears on a question,” for then we can distinguish between kinds of reasons “by distinguishing between kinds of questions on which a consideration can bear.” (Hieronymi, “The Wrong Kind of Reason,” 437–38.) A consideration becomes a reason when it stands in a certain relation to the question on which it bears. That is, it becomes a reason because of the relation “between the question on which the consideration bears and the attitudes of which it counts in favor” (*ibid.*, 438). Hieronymi suggests that we consider the relation between “*settling a question and forming or revising an attitude*” (*ibid.*, 447).

appropriate are reasons of the right kind, reasons that bear on the question of whether or not the object has the evaluative features that the attitude presents it as having.

D'Arms and Jacobson argue that we commit what they call the "moralistic fallacy" when we infer that an emotion is unfitting just because it would be morally wrong to feel that emotion.⁴⁰ Similarly, we might fallaciously infer that an emotion is unfitting because it would be harmful, inexpedient, or otherwise imprudent to feel the emotion. To illustrate the problematic conflation of fittingness with other criteria, D'Arms and Jacobson offer several examples: for instance, a joke may have the evaluative feature of being funny, or a colleague's promotion may have the evaluative feature of being enviable. The emotion of amusement presents its object – a joke – as being funny; the emotion of envy presents its object – a colleague's promotion – as being enviable. There are circumstances in which it would be morally wrong or inexpedient to feel amusement or envy – such as when a funny joke is racist or when you are depending on your enviably promoted colleague for an affirmative vote on your own personnel case – but the moral wrongness or imprudence of feeling a certain emotion may have no bearing on whether the emotion is fitting. A tendency to moralize everything makes us reluctant to see that an emotion that is morally wrong to feel could still be fitting; because we know that it is wrong to be amused by a racist joke, we do not like to admit that such a joke could still be funny. We may resist the point less when we are asked to distinguish between the imprudence of feeling an emotion and the fittingness of the emotion. For instance, it is quite plausible that while there are prudential reasons not to

When the relation is a *constitutive* relation, then settling the question "amounts to forming the attitude" (ibid.). For instance, settling the question of "whether *p*" amounts to forming the attitude of belief (that is, the belief that *p*), so "reasons that (are taken to) bear on whether *p* [are] the 'constitutive' reasons for believing *p*" (ibid.). We believe *p* by virtue of settling the question of whether *p*, and we settle the question when we are convinced by the reasons that (are taken to) bear on the question of "whether *p*." So too with evaluative responses such as resenting. Settling the question (in the affirmative) of whether you slighted me (given that slighting matters to me) amounts to forming the attitude of (for instance) resentment, so reasons that are taken to bear on whether you slighted me (given that slighting matters to me) are constitutive reasons for resenting you. In contrast to *constitutive reasons* for an attitude, when one develops an attitude for *extrinsic reasons*, finding the reasons convincing is not constitutive of having the attitude.

The fact that resenting you will teach you a lesson is (even given that I value teaching you a lesson) an extrinsic reason for resenting you; by being convinced that resenting you will teach you a lesson, I do not thereby form the attitude of resentment. Settling the question of whether resenting you will teach you a lesson does not have a constitutive relation to having the attitude of resentment. It merely motivates me (if I want to teach you a lesson) to try to cultivate the attitude of resentment. Importantly, it tells me nothing about whether your quality of regard has certain evaluative features, namely whether it is a form of disregard, a slighting.

⁴⁰ D'Arms and Jacobson, "The Moralistic Fallacy."

fear the wolf that is coming toward you, since the wolf will smell your fear, you nevertheless cannot infer from the fact that it would be unwise to fear the wolf that the wolf is not fearsome; it *is* fearsome, and fear is a fitting, and in that sense appropriate, response to the wolf.⁴¹

Let us consider, then, what concerns might lead someone to say that it is inappropriate for a person to hold themselves responsible for an unavoidable action by having a self-reactive attitude that presents the action as a wrongdoing. The first objection might be based on a concern about the consequences of holding oneself responsible in this way: when veterans hold themselves responsible for unavoidable but terrible actions, as they commonly do in cases of moral injury, this has seriously harmful effects on them. Recall that moral injury often manifests not just in feelings of guilt or shame, but also in self-destructive behaviors such as drug and alcohol abuse, self-harm, and (attempts at) suicide. Hence, one might correctly conclude that there is a reason for veterans *not* to hold themselves responsible even for terrible actions that they committed in the course of their service: doing so is self-destructive. This reason, however, would be an *extrinsic* reason: it does not bear on the question of whether or not their actions have certain evaluative features, such as the feature of being a wrongdoing, or being unthinkable, in ways that matter to them. The self-reactive attitude of guilt could fittingly reflect the way that it matters to Nik, for instance, that he killed a child; his feeling that “it’s wrong”⁴² might fittingly present that killing as a wrongdoing for which he holds himself responsible. The good reason that he has to *not* experience this self-reactive attitude (namely, whatever harm or distress the feelings or associated behavior may cause him) does not bear on the question of whether or not the killing has the evaluative features that the self-reactive attitude presents it as having.

Another concern that might underlie the charge that it is inappropriate for someone to hold themselves responsible for an unavoidable action is a concern about fairness or a sense that veterans do not deserve to be the targets of the self-reactive attitudes by which they hold themselves responsible. This objection, too, conflates questions of fit with other kinds of concerns. Whether or not someone *deserves* blame is a consideration that bears on the question of whether they should be punished, but not on the question of whether their reactive attitude (which may indeed be a blaming attitude) correctly reflects the evaluative features of the object, such as the feature of being a wrongdoing. What the evaluative features of the object are depends on what rightly

⁴¹ Ibid., 87.

⁴² Wood, *What Have We Done*, pp. 13–14.

matters to the person whose reactive attitude it is. As Pamela Hieronymi suggests, “mattering, not meriting, is... central” to moral responsibility.⁴³ What she calls the “merited-consequences conception” of moral responsibility – according to which a reactive attitude is fitting if it is the reaction that someone’s moral failure deserves or merits, the way one can be deserving of punishment – rests on the wrong assumption that reactive attitudes are voluntary in the way that actions are voluntary. The reactive attitudes “are, like belief, attitudes for which we can be asked our reasons, but which we cannot adopt at will.”⁴⁴ If they were voluntary and could be adopted at will (the way one can choose at will to punish or not punish someone), they could be directly responsive to extrinsic reasons, including reasons of merit. Instead, reactive attitudes are nonvoluntary responses that reflect what matters to us by presenting the object of the reaction (such as our own act of killing someone) as having a certain (dis)value (to us). It is what matters to the person who has the self-reactive attitude that provides reasons of the right kind – constitutive reasons – for the self-reactive attitude; these reasons are considerations that bear on the question of what matters about the action, but it is important to note that this does not necessarily mean that they are *good* reasons. Someone might be wrong about what matters; they could be making a mistake when their self-reactive attitude presents an action as having certain evaluative features; in that case, we could say they have *bad* reasons of the right kind. In the next section, I will consider whether fairness *must* matter to someone in order for them to not be making a mistake about what matters, that is, in order for their reactive attitude to reflect what *rightly* matters. So far, I have simply pointed out that the question of what rightly matters determines the *fittingness* of the reactive attitude, and this is a separate question from the question of whether there are extrinsic reasons for or against having a reactive attitude of blaming someone.

The Unthinkable

What rightly matters to service members and veterans about their own terrible but unavoidable actions? I am posing these questions in order to determine whether the distress expressed by morally injured veterans fits its

⁴³ Pamela Hieronymi, “I’ll Bet You Think this Blame is About You,” in Justin Coates and Neal Tognazzini (eds.), *Oxford Studies in Agency and Responsibility Vol. 5: Themes from the Philosophy of Gary Watson* (Oxford: Oxford University Press, 2019), 60.

⁴⁴ *Ibid.*, 66. In this way, a reactive attitude such as resentment is not *about* the one who is being resented, in the sense that it is not about whether or not they deserve the resentment; rather, it is about the one who resents, in the sense that it is about what matters to them.

object, namely whether it reflects what rightly matters to them, regardless of whatever extrinsic reasons there may be for them to not be distressed. Another way to put this is to ask what kinds of normative expectations of themselves they may rightly hold themselves responsible for meeting. Putting it this way may help us determine whether what people normatively expect of themselves can justifiably differ from what other people normatively expect of them.

What I want to suggest is that in some situations, it is fitting to have a self-reactive attitude that presents our treatment of others – for instance, killing someone – as *unthinkable* (in a sense to be elaborated). Morally injured veterans who hold themselves responsible for unavoidable acts such as killing may be expressing exactly this: that taking a human life remains unthinkable, even when it was the best thing they could have done. They may normatively expect themselves to avoid doing the unthinkable, even when the unthinkable is unavoidable, though this is not something that other people may normatively expect of them and thus hold them responsible for.

Strawson characterizes what matters to us in interpersonal relationships as goodwill, or at least lack of ill will, and other theorists following Strawson have focused on empathy.⁴⁵ But how we expect normal empathy to manifest in the way people treat each other depends on the level of vulnerability of one person to another. An empathic person who holds another's life in their hands, and who does not dehumanize the other person in any way, will generally find it to be *unthinkable* to kill the other person; in more intimate interpersonal relationships (perhaps between fellow soldiers rather than between a soldier and an enemy) someone may find it unthinkable to sacrifice the other person's life by failing to save them or letting them die. Faced with a situation in which we might have to kill or leave someone to die, we may have the experience of what Harry Frankfurt

⁴⁵ Other theorists working on the topic of reactive attitudes have developed the notion of quality of will in detail and have considered what qualities of will we demand from others. For instance, Shoemaker's tripartite division of senses of responsibility – attributability, answerability, and accountability – is organized so that each sense of responsibility corresponds to something different that we care about in the quality of other people's wills: their character, their judgment, or their regard; Shoemaker, *Responsibility from the Margins*, p. 24. Different reactive attitudes are fitting responses to the three different "objects": it is fitting to feel admiration/disdain in response to someone's quality of character, approval/disapproval in response to someone's quality of judgment, and gratitude/anger in response to someone's quality of regard; *Ibid.*, p. 26. For an agent to be an appropriate target of anger, the agent must be capable of the quality of regard that we seek in our interpersonal relationships, which Shoemaker describes as the complicated form of empathy that is necessary for the agent to avoid slighting others.

labels “volitional necessity,” in which we cannot bring ourselves to will something because we find it to be unthinkable.⁴⁶

Obviously, this response – finding it to be unthinkable to kill another person – can be overcome. Military training would be ineffective if it could not overcome the prohibition against killing. But, as the phenomenon of moral injury is making evident, many people who overcome the prohibition in the sense that they do manage to kill do not overcome the feeling that they are doing – or, if it is in retrospect, have done – something unthinkable. What so many of the morally injured veterans express is that they consider themselves to have transgressed a moral requirement, even though they “had no choice” – this *is* an expression of the unthinkability of their action, for to find it to be unthinkable to do something is different from just finding it to be wrong to do something: when something is unthinkable, we take ourselves to be morally prohibited from doing it even if it is impossible not to do it.

Just as a joke is funny whenever amusement is a fitting response to it, a colleague’s promotion enviable whenever envy is a fitting response to it, and a wolf fearsome whenever fear is a fitting response to it, an action can be *unthinkable* whenever the experience of volitional necessity is a fitting response to it, namely when someone’s having this response reflects part of what rightly matters to them. Unthinkability can be thought of as an evaluative feature of certain actions. The fitting response consists of taking ourselves to be nonnegotiablely prohibited from performing the unthinkable action; put differently, in experiencing volitional necessity we find certain actions to be impossible for us to will. Since we experience it as impossible – *in the sense of unthinkable* – to will these actions, *not* performing them is required even if it is also impossible, in the more ordinary sense of “impossible.” When we take an *unavoidable* action to be unthinkable, then when we unavoidably perform the action (as service members do in

⁴⁶ Harry Frankfurt, “The Importance of What We Care About,” in *The Importance of What We Care About* (New York: Cambridge University Press, 1988), 86. This is not only true of civilians who never encounter a situation in which they might have to kill someone, but also true of combat soldiers – it is part of what training must overcome, but training may be only partially successful in this respect. Dave Grossman’s *On Killing* reports in great detail on the psychological resistance to killing that soldiers experience and explains how military training can and does function to overcome this resistance; Dave Grossman, *On Killing: The Psychological Cost of Learning to Kill in War and Society* (New York: Back Bay Books, 1995). Although Grossman’s work draws in part on a study by S.L.A. Marshall that has been called into question (particularly his claims about the numbers of soldiers who could not bring themselves to fire their weapons to kill) Grossman’s main points seem to hold. For an interesting study of people’s (civilians’) “aversion to harmful action,” see Fiery Cushman, et al., “Simulating Murder: The Aversion to Harmful Action,” *Emotion*, 12, no. 1 (2012), 2–7.

situations, for instance, in which killing is the best thing they can do), we must take ourselves to have failed to meet the impossible requirement not to perform the action. The recognition of this failure is what is reflected in the distressed emotions of moral injury. The distress is a fitting self-reactive attitude because it correctly apprehends the evaluative features of our having done what we unavoidably did – namely the unthinkability of it – and with it, our own failure. Even if we had no control over performing the action, we have failed to meet the impossible requirement not to perform it. The anguished sense of responsibility characteristic of moral injury reflects the recognition of the failure to meet the normative expectation not to do the unthinkable, even if it is unavoidable.

The normative expectation that someone avoid the unavoidable, however, is an unfair expectation. Is this relevant? I believe that, in general, unfairness matters when we apply the normative expectation to other people by reflecting it in our interpersonal reactive attitudes, but not necessarily when we apply it to ourselves by reflecting it in our self-reactive attitudes. In short, we are permitted to have unfair expectations of ourselves. This is because what rightly matters to us – what we can hold ourselves responsible for – includes two kinds of values: values that we share with other members of our society, and solely subjective values. We may have a subjective experience of valuing that is not tempered or limited by concerns about fairness, and this may be the basis for the normative expectation that one avoid doing something unthinkable even if it is impossible to avoid. Subjective experiences of valuing might also lead one to have expectations of oneself for levels of self-sacrifice that others would regard as supererogatory and that others could not justifiably expect of one.⁴⁷

In contrast, normative expectations may only rightly come to be shared in certain ways, and an interpersonal reactive attitude is a demand for someone else to meet either what *are* shared normative expectations or what one, by virtue of having a particular interpersonal reactive attitude, is proposing *should* be a shared normative expectation. Norms of fairness, reciprocity, or justifiability are often cited – for instance, in contractualist

⁴⁷ Bernard Williams has made this point: “Ethically outstanding or possibly heroic actions... , in being more than obligations, are not obligatory, and we cannot usually be asked to do them or be blamed for not doing them. But the agent who does such a thing may feel that he must do it, that there is no alternative for him, while at the same time recognizing that it would not be a demand on others. The thought may come in the form that it is a demand on him, but not on others, because he is different from others; but the difference will then typically turn out to consist in the fact that he is someone who has this very conviction”; Bernard Williams, *Ethics and the Limits of Philosophy* (Cambridge, MA: Harvard University Press, 1985), p. 188–89.

moral theory – as regulating a hypothetical process through which further normative expectations can be constructed; one way to construe this is to say that people may not claim a normative principle ought to be shared if others, properly motivated by the aim of finding shared principles, could reasonably reject it.⁴⁸ Despite the fact that norms of fairness may in practice be violated, or that the concept of fairness can be interpreted in conflicting ways and does not necessarily imply equality, it is still the case that many actual societies have developed norms of fairness or reciprocity that play some role in regulating the construction of further norms.⁴⁹ Our interpersonal reactive attitudes must either reflect these fairly constructed normative expectations or serve to fairly propose that some normative expectation be shared.

We may hold ourselves responsible for more than what we may hold most other members of our society responsible for, because our expectations of ourselves may be based on a subjective experience of valuing that is not tied to such a process of norm construction, and thus the self-reactive

⁴⁸ T.M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1988). This does not mean that actual people do not try to hold others responsible in ways that reflect unfair or unjustifiable normative expectations or that actual people do not (perhaps in a self-deceived way) put forward claims that appear to be fair when they are not. Shared normative expectations may very well be unfair (often because they apply differently to members of different social groups); and people may very well have interpersonal reactive attitudes that reflect what they believe should be shared normative expectations, even though these expectations would be unfair. The common phenomenon of “blaming the victim” is based on unfair normative expectations; this is evident, for example, when a woman is held responsible for being raped because she was walking alone at night or dressed in a way that the rapist found to be provocative, or a black man is held responsible for being subject to excessive force during an arrest because he acted in a manner that caused a police officer to feel threatened.

⁴⁹ One can argue (on the basis of a claim about the function of morality, given certain assumptions about human nature and about what human cooperation requires) that any moral system that does not have some norm of fairness or reciprocity, and some recognition of the need for justification of shared normative expectations, is inadequate. David Wong, in developing a constructivist position that he calls *pluralistic relativism*, argues that there are some universal constraints on any adequate morality, including that any adequate morality must have some norm of reciprocity and some requirement that norms be justifiable to those who are to be governed by them; David Wong, *Natural Moralities: A Defense of Pluralistic Relativism* (New York: Oxford University Press, 2006). He claims that “some form of reciprocity is a norm for all cultures we know, where reciprocity is conceived as a fitting and proportional return of good for good” (*ibid.*, p. 47), noting that, “while reciprocity in a very general sense is the norm for human cultures, the specific form it takes varies a great deal across cultures”; *ibid.*, p. 50. Regarding the constraint of justifiability, he points out that “as a system for promoting cooperation, morality works through a large degree of voluntary acceptance of its norms... If conformance to its norms... depended solely on the threat of force or coercion, the costs would detract greatly from the benefits of social cooperation itself... voluntary acceptance of moral norms came to be seen as based on their justifiability to those governed by them. Hence another constraint on moralities is that justification for following the norms and reasons of an adequate morality cannot crucially depend on falsehoods”; *ibid.*, p. 59.

attitudes characteristic of moral injury can be fitting. I may not justifiably impose on others expectations based on what matters only to me, that is, based on values I could not expect them to accept, but I may hold myself to these expectations, for I *have* in fact accepted them. When, through self-reactive attitudes, I hold myself responsible specifically in relation to my subjective values, the fact that I hold *only* myself responsible means that there is no chance of holding someone responsible for values that they neither share nor could be reasonably expected to share. For shared values, the question to be settled is not what matters to *me*, but what matters to *us*, and if determining what matters to us entails knowing what we can reasonably expect each other to accept, shared normative expectations will be limited in ways that my expectations of myself are not.⁵⁰ While I can certainly hold others responsible for not killing when killing is an action that they *can* avoid, I cannot reasonably expect others to avoid the unavoidable or to emerge from a moral dilemma without committing an action that would otherwise be prohibited. But through the subjective experience of volitional necessity, I may very well require these sorts of impossible things of myself.

Conclusion

How do we answer practical questions about how to respond to the problem of moral injury and its associated distress, if, as I have argued, the distress is fitting, and yet there are extrinsic reasons that recommend against it? One part of the answer is obvious: we should undoubtedly work to reduce the *causes* of moral injury, whenever they can be reduced. Moral injuries experienced by perpetrators of *avoidable* wrongdoings can be prevented to the extent that the perpetrators can be prevented from committing the wrongdoings. Even regarding wrongdoings that were unavoidable *given the situation someone was in*, it may be that their being in that situation in the first place was completely avoidable: bad political decisions (including the decision to go to war), inadequate efforts at diplomacy, and poor military leadership can all result in service members being put in situations in which they have no better options than to commit moral wrongdoings. Progress in democratic politics, better diplomacy, and improvements in

⁵⁰ Scanlon's contractualism leaves him room to accept what I am claiming here, because he acknowledges that the contractualist theory applies only to one circumscribed area of morality – the realm of “what we owe to each other” – and that morality more broadly understood admits a greater plurality of values; Scanlon, *What We Owe to Each Other*.

military leadership could prevent some moral injuries.⁵¹ Nevertheless, it is a sad truth that not all of the causes of moral injury can be eliminated, because moral failure is not always avoidable: there will always be moral luck, epistemic limitations, and conflicts between incommensurable values, that lead in different ways to unavoidable failures, as well as situations that force people to choose between wrongdoing and self-sacrifice. I am not a pacifist and do not believe either that war can always be avoided or that war can be conducted without warriors having to commit acts that it is fitting to treat as unthinkable.

A separate question concerns how others should respond to a veteran who expresses distress once a moral injury has taken place and the veteran is suffering from it. In the case of veterans whose moral injury stems from a wrongdoing that they could have and should have avoided, both self-blame and blame from others are fitting and serve to hold the person responsible for meeting fair normative expectations. Here the therapeutic approach called Adaptive Disclosure makes sense conceptually.⁵² This approach involves the therapist affirming the appropriateness of the veteran holding themselves responsible for their wrongdoing: it “allows for accurate and legitimate self-blame when warranted but also promotes the possibilities of self-forgiveness, compassion, and moral reparation”,⁵³ and “encourages a recommitment to pre-event personal ethical and moral standards.”⁵⁴

Responding well to a veteran whose moral injury stems from their experience of *unavoidable* moral failure, on the other hand, depends on accepting the fittingness of the veteran’s attitude of holding themselves responsible for the wrongdoing, while meanwhile recognizing that there are extrinsic reasons for diminishing the distress. The challenge is to reduce the suffering *without* engaging in any denial of the fittingness of the self-reactive

⁵¹ Preventing moral injury by improving military leadership is one of the main aims of Jonathan Shay’s work; Shay, *Achilles in Vietnam; Odysseus in America*; and “Moral Injury.” While he employs a definition of “moral injury” that is focused on how one can be injured by being betrayed by people who are in a position of authority, he also recognizes that such betrayals, or failures of leadership more generally, can put service members in positions in which they themselves must commit wrongdoings.

⁵² See Litz, et al., “Moral Injury”; Shira Maguen and Brett Litz, “Moral Injury in Veterans of War,” *PTSD Research Quarterly*, National Center for Posttraumatic Stress Disorder, 23, no. 1 (2012), 1–3; and Gray, Nash and Litz “When Self-Blame is Rational and Appropriate.” While I endorse the conceptual assumptions behind this approach, I am not in a position to evaluate the approach empirically or to say anything about its therapeutic effectiveness.

⁵³ Gray, Nash and Litz “When Self-Blame is Rational and Appropriate,” 383.

⁵⁴ *Ibid.*, 386.

attitude, and thus to leave room for them to hold themselves responsible and to express that they have failed morally.⁵⁵ In insisting that some of the actions that it may be beyond our control to avoid are truly unthinkable, morally injured veterans are actually offering anyone who will listen to them an important but tragic truth that we should treasure even if we would prefer the fantasy that moral goodness is always possible.

⁵⁵ Bryan Doerries' "Theater of War," I believe, may achieve exactly this kind of response to morally injured service members and veterans. See <https://theaterofwar.com/>.