


RESEARCH ARTICLE

# Relational Fault and Unforeseeable Victims

Adam Slavny 

School of Law, University of Warwick, Coventry CV4 7AL, UK  
Email: a.slavny@warwick.ac.uk

(Received 21 February 2025; Revised 19 March 2025; Accepted 28 February 2025)

## Abstract

In this paper, I argue against a widely held view about interpersonal moral relations inspired by Benjamin Cardozo's landmark judgment in *Palsgraf v Long Island Railroad Company*, which I call the *Relational Fault Requirement*. The requirement holds that in order for A to commit a directed wrong against B, A must be at fault in relation to B. I present two ways of understanding wrongs that violate this requirement: (1) that one is wronged if one is harmed by a wrongful action, and (2) that one is wronged if one is harmed by a wrongful action *and* the outcome one suffers is sufficiently similar to the grounds on which the action is wrong. Accepting either of these ideas requires rejecting the *Relational Fault Requirement* and encourages us to rethink the core elements of directed wrongdoing.

## I. Introduction

A distinction is often made between acting wrongly and wronging another. To say that Ali wrongs Cat, for example, is to say more than that Ali did something he should not have done. It means Ali's act bears some moral relation to Cat, and usually, this is reflected in the fact that Cat is entitled to make certain demands of Ali, such as explanation, apology, and compensation.

Directed wrongs appear to have what I will call a *Relational Fault Requirement*:

*Relational Fault Requirement*: In order for A to commit a directed wrong against B, A must be at fault in relation to B.

Perhaps not all directed wrongs have this requirement because not all are fault-based. For example, if promise-breaking is rightly regarded as a strict liability wrong, Ali wrongs Bea if he fails to fulfill his promise to her, regardless of whether he is at fault for doing so. Nevertheless, the requirement plausibly applies to a range of fault-based wrongs. I will primarily discuss the wrong of negligent harming, although the relevant views apply to other fault-based wrongs such as reckless and intentional harming.

The *Relational Fault Requirement* takes inspiration from Benjamin Cardozo's landmark judgment in *Palsgraf v Long Island Railroad Company*.<sup>1</sup> This case concerned an employee of the defendant railroad company who negligently pushed a passenger onto a moving train. The passenger dropped a package that happened to contain fireworks, which exploded and caused a tall scale standing on another part of the platform to fall onto the plaintiff. Cardozo, ruling for the defendant, held that the plaintiff was not owed a duty of care because she did not belong to that class of people foreseeably endangered by the defendant's negligent act. The defendant's negligence may have violated a duty owed to the passenger, but not a duty owed to Palsgraf.<sup>2</sup>

Although the question before the New York Court of Appeals in this case concerned a claim for the tort of negligence, the *Relational Fault Requirement* arguably reflects a deep feature of our interpersonal moral relations. Moreover, while the point at issue in Palsgraf was the existence of a duty of care (rather than the breach of that duty) it is natural to extend the kind of relationality embodied in the decision to fault. Those directed moral wrongs that require some degree of fault also require that fault to have the right relation to the wronged party.

The *Relational Fault Requirement* does not insist that A must be at fault exclusively in relation to B. A can be at fault in relation to many people, as in the case of a traffic pile-up or large-scale accident, while being at fault toward each of the individual victims. But it does rule out two other forms of fault. The first is a fault that is not directed at anyone at all—so-called “negligence in the air.” If A is at fault *tout court*, without being at fault in relation to B, A does not wrong B in the directed sense. The second type of fault ruled out is fault directed at a third party. In Palsgraf, the defendant was at fault in relation to the passenger, but the plaintiff could not piggyback on this to derive an action in negligence. To claim successfully, she had to show that he was at fault in relation to her, in the form of a breach of a duty of care owed to her.

The *Relational Fault Requirement* is intuitive, but it raises some puzzles. We can see this via the following trio of cases. First, consider:

*Quarry*: Ali sets up an explosion in a quarry. Hikers often pass through the area, and he fails to erect a fence that would keep them safe. Bea is close by, and she follows a hiking trail into the quarry. Bea is injured in the explosion.

*Quarry* is a standard negligence case. Ali's failure to erect a fence is negligent with respect to hikers such as Bea. Next, consider:

*Quarry 2*: Ali sets up an explosion in a quarry. No one ever passes through the area as it is far from any hiking trails. Erecting a fence would prevent paragliders from landing in the quarry, however, no one is permitted to paraglide in this area so it is unforeseeable that paragliders would be at risk. Ali fails to erect the fence. Cat, who is paragliding miles away where the activity is permitted, is blown into the quarry by an unforeseeable storm. Cat is injured in the explosion.

<sup>1</sup>248 N.Y. 339 (1928).

<sup>2</sup>There have been similar decisions in English negligence cases too. In *Bourhill v Young* [1943] AC 92, for example, a woman who miscarried as a result of hearing a motorcycle crash and later seeing blood on the road was not able to recover damages from the person who negligently caused the crash. The defendant owed a duty not to drive negligently, but that duty was owed to other road users rather than the claimant.

In *Quarry 2*, Ali does not wrong Cat. His failure to erect the fence is not at fault in relation to them, since it is unforeseeable that they would be endangered. It is also not at fault in relation to anyone else, since Ali has no reason to think anyone will be in the vicinity of the explosion. Some might think Ali still owes Cat an apology or compensation. If so, this would have to be based on some other ground, such as the fact that Ali causes harm to Cat, or that Ali is outcome responsible for the harm to Cat, or that Ali is the primary beneficiary of the act that harms Cat, or some such.

Finally, consider:

*Quarry 3*: Ali sets up an explosion in a quarry. Hikers often pass through the area, and he fails to erect a fence that would keep them safe. Bea is one such hiker whom Ali puts at mortal risk, although fortunately, she avoids the explosion. Cat, who is paragliding miles away, is blown into the area by an unforeseeable storm and lands in the quarry. Cat is injured in the explosion.<sup>3</sup>

What we ought to say about *Quarry 3* is more of a puzzle. On the one hand, Ali's relation to Cat seems the same as it is in *Quarry 2*. It is unforeseeable that Cat will be injured as they are paragliding miles away. So, if we focus only on Ali's relation to Cat, it seems that Ali does not wrong them.<sup>4</sup> On the other hand, Ali could and should have foreseen harm to others such as Bea and thus he is at fault for not erecting the fence. Had he done what he morally ought to have done, he would not have harmed Cat. We might think this distinguishes his conduct in *Quarry 3* from that in *Quarry 2*, where his actions are not wrong in relation to anyone.

If we adhere to the *Relational Fault Requirement*, we should conclude that Ali does not wrong Cat in *Quarry 3*. The fact that his conduct is at fault in relation to hikers such as Bea is not sufficient. I will argue that we should reject this claim. We should accept what I take to be the intuitive verdict that Ali wrongs Cat in *Quarry 3* as he harms them through his negligent failure to take care. The wrong Ali does to Cat can piggyback on Ali's fault in relation to hikers such as Bea. Consequently, we must revise the *Relational Fault Requirement*, perhaps quite radically. Not only that but cases like *Quarry 3* show that a person can be wronged even when harm to them is unforeseeable. I will suggest that we can nevertheless vindicate the idea of directed wrongdoing in these cases, and then address some objections that may be raised against this more expansive account.

## II. Wrongs to Unforeseeable Victims

I suggested that the *Relational Fault Requirement*, which takes its inspiration from *Palsgraf*, implies that Ali does not wrong Cat in *Quarry 3*. However, in a recent paper, Jed Lewinsohn offers an analysis of Cardozo's judgment in *Palsgraf* that, if extended to cases like *Quarry 3*, has a different implication. Lewinsohn interprets Cardozo as articulating a principle he calls *Danger Zone*, which holds that, if an agent believes or should believe that some act imposes an unreasonable risk of harming someone or

<sup>3</sup>For a discussion of similar cases, along with a number of variations, see Jed Lewinsohn, "I Didn't Know It Was You": *The Impersonal Grounds of Relational Normativity*. 59 *Noûs* 191 (2025), hereafter "I Didn't Know It Was You." We will return to Lewinsohn's arguments shortly.

<sup>4</sup>R. Jay Wallace incorporates a similar constraint into his contractualist theory. See *The Moral Nexus*, (Princeton University Press, 2019) Ch. 5.

other who instantiates a “danger-making property,” then the agent owes a duty of care to all and only those who even unforeseeably instantiate that property.<sup>5</sup> A danger-making property is roughly one that carries a relevant risk when that property is instantiated, such as being in the physical proximity of an explosion or drinking from a contaminated water source. A person who instantiates this property is owed a duty of care and may be wronged if this duty is breached, even if it was unforeseeable that this person would instantiate the property.

Importantly, this principle generates duties via these danger-making properties rather than the ex ante probability of harm. The ex ante probability of Cat being hurt by the explosion is extremely low since they are paragliding miles away, but once Cat is in the quarry, they are owed a duty because they instantiate a danger-making property. Similarly, Mrs. Palsgraf would have been able to claim against the defendant if, say, the box had fallen on her feet, regardless of whether the ex ante probability of her being in the vicinity of the box was very low. If this is right, and if we extend this idea to the realm of moral duties, the *Palsgraf* decision provides support for the intuition that Cat is wronged in *Quarry 3*.

Lewinsohn’s primary interest is in defining which victims of a wrong are “properly aggrieved,” meaning that they possess a set of entitlements in relation to the wrong such as the right to demand an apology, compensation, and so forth. Directed wrongdoing and proper aggrievement are separate concepts, but since I think they typically go together, for my purposes I will treat *Danger Zone* as applying to both. Lewinsohn appeals to his analysis to argue against a range of theories that seek to ground proper aggrievement in some relation between the wrongdoer’s act and that individual’s interests. A general theme of the argument is that none of these interest-based theories can accommodate *Danger Zone* (and thus *Palsgraf*) because an appeal to the ex ante probability of harming any particular person cannot predict who will be properly aggrieved. Two people with equal ex ante probabilities of harm might find themselves on different sides of the danger zone boundary. This would happen if, say, Bea and another hiker, Delta, who are equally at risk, choose different paths, only one of which leads into the quarry. His conclusion is that the grounds of proper aggrievement are mediated through impersonal danger-making properties and cannot be accounted for using the “highly individualistic notions that have dominated philosophical discussions of relational normativity.”<sup>6</sup> The best explanation for how we demarcate the class of the properly aggrieved is a conventionalist one, which depends ultimately on a set of social and institutional practices.

I agree with *Danger Zone*’s implication that Cat is wronged in *Quarry 3* (and is properly aggrieved). While it is unforeseeable that Cat would be in the vicinity of the explosion, Ali’s wrongful failure to take care means that he wrongs anyone who is in fact in the vicinity. But this judgment about *Quarry 3* violates the *Relational Fault Requirement*. Had Ali not been at fault in relation to anyone, the case would be analogous to *Quarry 2* and Cat would not have been wronged. It is only because Ali’s failure to erect the fence is wrong in relation to hikers like Bea that we have the intuition that Cat is also wronged.

We could try and get around this point by saying the following. Ali is at fault in relation to Cat because he did something inherently dangerous (setting an explosive)

<sup>5</sup>Lewinsohn, “I Didn’t Know It Was You,” 194.

<sup>6</sup>Lewinsohn, “I Didn’t Know It Was You,” 192.

which in fact harmed Cat. It does not matter that it was not foreseeable that Cat would be in the quarry; all that matters is that it was foreseeable that Ali would harm Cat *if* they turned up in the quarry because of the dangerousness of his conduct. This is sufficient to conclude that Ali is at fault in relation to Cat. The problem with this suggestion is that it defines relational fault too broadly (if we were talking about the duty of care in negligence, we could say it defines this too broadly for similar reasons). It is true in both *Quarry 2* and *Quarry 3* that Ali would harm Cat *if* they turned up in the quarry, yet in the former case Ali clearly commits no wrong. In fact, the highest degree of care would not guarantee that Ali would not harm Cat *if* they turned up in the quarry, given the potentially harmful nature of explosives. We cannot appeal directly to the inherent dangerousness of Ali's conduct to explain why he wrongs Cat, because inherently dangerous conduct is only wrongful if it imposes risks on others. Detonating a bomb on Mars is not faulty in relation to anyone because no one is subject to an unjustified risk. On the other hand, if by inherently dangerous we mean to imply that the activity poses an unjustified risk to someone, we are back to the claim that Ali's wrong against Cat in *Quarry 3* piggybacks on his fault in relation to foreseeable victims like Bea. Ali's act of setting an explosive and failing to erect a fence imposes an unjustified risk on hikers, not paragliders. The conclusion that Ali nevertheless wrongs Cat must, therefore, violate the *Relational Fault Requirement*.

The second claim I want to make in relation to *Quarry 3*, and contra Lewinsohn, is that the wrong Ali commits against Cat is a nonconventionalist feature of relational normativity. I will explore two ways of understanding this wrong and conclude that it is the *Relational Fault Requirement* that needs to be revised rather than the notion that the wrong is nonconventional.

### III. Harming Through a Wrongful Act

To achieve this aim, we might endorse something like the following:

*Nonrelational Fault:* When A commits some act,  $\phi$ , it constitutes a directed wrong against B if (1)  $\phi$  is either (i) at fault in relation to a third party or (ii) at fault tout court and (2)  $\phi$  causes harm to B.

Here I understand fault to encompass negligent, reckless, and intentional wrongdoing. In *Quarry 3*, Ali's failure to erect the fence is negligent, but if it had been reckless or intentional this would also satisfy *Nonrelational Fault*. Often these faulty acts will be wrongful in themselves, especially when they impose unjustified risks on others as Ali does in *Quarry 3* (this assumes that the imposition of pure risks can be wrongful, which is a view I have defended elsewhere),<sup>7</sup> or when they constitute inchoate wrongs. I leave open the possibility that an act can be at fault in one of these ways without amounting to wrongdoing. However, I will focus here on acts that themselves constitute moral wrongs, whether to others or no one in particular.

Another important clarification is that, since *Nonrelational Fault* is a view about moral rather than legal wrongs, the wrongdoer must be negligent (or reckless or intentional) only in the moral sense, which may differ from how the law defines these concepts. For example, the objective standard of care enshrined in the law of negligence means that some people who are not morally at fault, because they lack

<sup>7</sup> Adam Slavny and Tom Parr, *What's Wrong with Risk?* 8(2) THOUGHT 76 (2019).

the capacity to live up to this standard, are nevertheless legally liable, and I do not intend to include these people in the scope of *Nonrelational Fault*.<sup>8</sup>

*Nonrelational Fault* sits alongside relational fault as a distinct route to directed wrongdoing. A directed wrong of the sort specified by *Nonrelational Fault* can piggyback on two kinds of nonrelational faulty acts: those that are at fault with respect to a third party and those that are at fault tout court. I will focus on acts like those in *Quarry 3* and *Palsgraf* that are at fault with respect to a third party and will mention cases where acts are at fault to no one in particular where relevant (note that I consider *Quarry 3* and negligence cases in general instances of relational fault even though we often do not know the identity of the foreseeable victims, or precisely how many there are. As long as there are some foreseeable victims and the risk imposition is unjustified, the act is faulty in relation to them). Like the *Relational Fault Requirement*, I do not intend *Nonrelational Fault* as a universally sufficient condition for directed wrongdoing. It is a sufficient condition only for those wrongs that have some fault requirement, especially negligent, reckless, or intentional harming.

With these clarifications out of the way, *Nonrelational Fault* implies that Ali wrongs Cat in *Quarry 3*. More generally, it implies that a person is always wronged if they are harmed by an act that is at fault in relation to a third party. One objection to *Nonrelational Fault* is that, in *Quarry 2* and *Quarry 3*, the relation between Ali and Cat is exactly the same. In both cases, it is not foreseeable that Cat will be in the quarry because they are paragliding miles away. How, then, can it be true that Ali wrongs Cat in *Quarry 3* but not in *Quarry 2*?

It is true that Ali harms Cat to the same degree in *Quarry 2* and *Quarry 3*, and Cat's ex ante probability of harm is also the same. So if we think directed wrongdoing is determined exclusively by the actual or expected harm imposed on the victim, this objection is powerful. Some views seem vulnerable to this objection, as Lewinsohn shows.<sup>9</sup> For example, in some contractualist views, a person's objection to a given principle of action is determined only by its impact on their personal interests, where these are understood in terms of actual or expected harm. On such views, if Ali does not wrong Cat in *Quarry 2*, we must reach the same conclusion in *Quarry 3*.<sup>10</sup>

There are two responses to this problem that offer alternatives to Lewinsohn's suggestion that the explanation is a matter of convention. One is to reject the narrow conception of personal interests, and another is to argue that facts independent of interference with personal interests sometimes ground directed wrongdoing. I prefer the second solution, but both are more plausible than adhering to the narrow conception of personal interests. Certainly, we cannot infer that Cat's interests are a setback to the same degree in *Quarry 2* and *Quarry 3* from the fact that Ali's relation to Cat is the same in the two cases because this is not true. In *Quarry 3*, Ali harms Cat through a negligent act, and in *Quarry 2* he does not. It is true that, in *Quarry 3*, Ali's act is negligent because of the risk it imposes on others, but this does not change the fact that this relation—that of harming through a negligent act—holds specifically between Ali and Cat.

<sup>8</sup>For an argument that this class of defendants undermines nonpragmatic justifications of the objective standard, see Adam Slavny, *Wrongs, Harms, and Compensation: Paying for Our Mistakes* (Oxford, 2023) Ch. 4.

<sup>9</sup>Lewinsohn, "I Didn't Know It Was You," especially 197–204.

<sup>10</sup>Lewinsohn's primary foil here is the version of contractualism developed by Wallace in *The Moral Nexus*, especially Ch. 5.

Even the “individualist restriction” adopted by contractualists and criticized by others<sup>11</sup> can be modified to accommodate this idea. If we expand the notion of personal interests to include whether one’s interests are impacted by an act that is independently negligent, reckless, or showing culpable intent, then Cat’s personal interests are more severely affected in *Quarry 3* than in *Quarry 2*. This is not such an unusual or arbitrary extension of the notion of personal interests. My primary concern with others’ conduct is that they do not harm me, but I also care that they do not harm me through a moral failing on their part, even if that failing is not directed at me. This is especially evident when the failing is highly culpable. All else equal,<sup>12</sup> I would rather have my arm broken by an honest accident than a botched attempt to kill the person standing next to me. In the latter case, my injury results from behavior that I, and all others, have strong reasons to condemn. It is something that morally should not have happened, and that should have been prevented by the person who caused it.

Of course, we need not spell this thought out in terms of personal interests. We could go with the second option and hold that being harmed by a faulty act plays an interest-independent role in grounding directed wrongs. This option will appeal more to those who, like me, adopt a pluralistic attitude toward the grounds of directed wrongdoing, rather than trying to reduce them to a single general idea. Having raised these options, I will remain ecumenical instead of trying to adjudicate them any further. Our next task is to ask, aside from explaining the intuitive difference between *Quarry 2* and *Quarry 3*, why we should think *Nonrelational Fault* is plausible.

#### IV. The Appropriateness of Guilt and the Inappropriateness of Mere Agent-Regret

Consider agent-regret. In Bernard Williams’ classic formulation of agent-regret, its “constitutive thought” is that *I* am the agent of harm, rather than that I have done something I morally should not have done.<sup>13</sup> In *Quarry 3*, I suggest it would be inappropriate for Ali to experience mere agent-regret, as this would treat his harming Cat as equivalent to his causing an innocent accident, like in *Quarry 2*. He is a culpable agent of harm, not an innocent one, even though the culpable wrong that brought about that harm is directed at a third party. The flip side of this claim is that it would not be inappropriate for Ali to feel full-blooded guilt toward Cat, since he did something morally wrong which hurt Cat. Similarly, Cat is entitled to be angry about the fact that they would not be in this situation if Ali had adhered to his moral obligations. This forms the basis of a distinctive complaint that is not available to the victim of an innocent accident. In *Quarry 2*, for instance, where Ali does nothing wrong and yet harms Cat, it would be unfair to expect Ali to feel full-blooded guilt rather than agent-regret.

Another way of seeing why mere agent-regret is inappropriate is that it would be unwarranted for third parties to talk Ali out of his negative feelings. While we expect those who cause innocent harm to feel agent-regret, we also think it is appropriate for

<sup>11</sup>See Joe Horton (2017) *Aggregation, Complaints, and Risk* 45 PHI & PUB. AFF. 54 (2017).

<sup>12</sup>This caveat is necessary as we might prefer to be culpably harmed for other reasons, such as having a greater prospect of recovering compensation. Thanks to Zoe Sinel for this point.

<sup>13</sup>Bernard Williams, *Moral Luck* (Cambridge, 1981).



others to try and lessen their pain by reassuring them that they were not at fault.<sup>14</sup> This is not true of guilt, where we would look askance at anyone who tried to talk a culpable wrongdoer out of feeling guilty. This difference is an important way to distinguish between agent-regret and guilt. In *Quarry 2*, it would be appropriate for a third party to try and diminish Ali's sense of agent-regret with sentiments such as, "It wasn't your fault," "You did nothing wrong," and "No one else would have done anything differently." In *Quarry 3*, where Ali wrongly disregards the safety of hikers like Bea, similar sentiments would be inappropriate. This is even more evident if we increase Ali's culpability. If he acts recklessly or intentionally in not erecting the fence, it seems even more inappropriate to talk him out of his guilt.

At this point, one might press the following objection. Granted, Ali ought to feel guilt rather than mere agent-regret in *Quarry 3*, but this is because he is at fault in relation to hikers like Bea. Either he should feel a directed form of guilt, but to the hikers rather than to Cat, or he should feel guilty in a nondirectional sense. In either case, the appropriateness of guilt and the inappropriateness of mere agent-regret can be explained without the assumption that Ali wrongs Cat.

The problem with this reply is that it mischaracterizes the phenomenology of the case. In *Quarry 3*, it would be odd if Cat were to confront Ali, nursing the life-changing injuries that he inflicted through his careless conduct, only to be told that he did not feel the least bit guilty toward Cat. Similarly, if Ali did feel guilty toward Cat, it seems inappropriate for third parties to contest this, saying, "Your action wasn't wrong in relation to *them*, so direct your guilt elsewhere!" To this, I think Ali could, and should, respond, "I just ruined this person's life because of my negligence, and you think I have no reason to feel guilty about that? Maybe we shouldn't hang out so much..."

To make the same point another way, consider the following cases:

*Two Responses 1:* Ali negligently harms Bea and innocently harms Delta. He feels bad for what he did to both, but he can only apologize to one of them.

*Two Responses 2:* Like *Quarry 3*, except Ali also innocently harms Delta. He feels bad for what he did to both, but he can only apologize to one of them.

In *Two Responses 1*, I take it that Ali should apologize to Bea rather than Delta. He has more reason to apologize for wrongfully harming Bea than for innocently harming Delta (if you do not think we ought to apologize for innocent harm at all, this gives you all the more reason to distinguish the responses owed to the two victims). I think *Two Responses 2* is analogous. Ali should apologize to Cat rather than Delta. To apologize to Delta instead, or flip a coin, would ignore the culpability that was implicated in the harm he did to Cat. This suggests that what he does to Cat is closer to a standard case of directed wrongdoing than it is to a case of innocent wrongdoing. His response goes beyond mere agent-regret and is therefore best understood as a form of guilt.

Appealing to nondirectional guilt to explain the phenomenology does not fare better. If Ali has reason to feel guilty only about the faultiness of his conduct, rather than because he wrongs anyone in particular, this seems like a self-indulgent form of guilt. There are contexts when preoccupation with one's own moral character is

<sup>14</sup>For a discussion of this feature of agent-regret, see Nichols and Kamtekar, *Agent-Regret and Accidental Agency* 43 *MIDWEST STUD. PHIL.* 181 (2019).



appropriate, such as when a person is engaging in moral self-examination. When it comes to guilt about a particular wrongful act, though, nondirectional guilt misses something important. If anything, it would be worse if Ali did not feel guilty in relation to Cat or Bea than if he were to feel it only in relation to Bea. Ali should feel guilty not just because he acted wrongly tout court, but because this wrong caused harm to another person. Ali's fault may not be directed at Cat, but once it causes harm to them, his guilt should be so directed.

## V. The Broadness Problem

Although I think *Nonrelational Fault* helps to explain how Ali wrongs Cat in *Quarry 3*, it faces challenges. The central problem is that it is very broad. Rather than giving us a way of distinguishing between *Quarry 3* and *Palsgraf*, it implies that the victims in both cases were wronged because they were harmed by an action that was at fault in relation to others. Yet it is more intuitive that Cat is wronged in *Quarry 3* than that the plaintiff was wronged in *Palsgraf*.

One response to the broadness problem is to appeal to the distinction between causation requirements and other elements of a wrong. One reason why cases like *Palsgraf* can be confusing is because, even if we agree with the result, it is not clear whether it should be based on Cardozo's rationale that the plaintiff was owed no duty, or on the idea that the relevant causal conditions were not satisfied. Consider the following case, which brings out these issues.

*Quarry 4*: Ali sets up an explosion in a quarry. He fails to erect a fence that would keep hikers safe. Bea is there, but by luck, she avoids the explosion, which goes off in an empty quarry. There is an unforeseeable structural weakness in a cave half a mile away, and the explosion triggers it to collapse. Owing to another unforeseeable structural weakness in a cliff another half mile away, the collapse triggers a landslide that injures Cat.

*Nonrelational Fault* implies that Ali wrongs Cat in *Quarry 4* because Ali's act is at fault in relation to Bea, and Cat ends up being harmed as a result. But it is less intuitive that Cat is wronged in *Quarry 4* than it is in *Quarry 3*. In *Quarry 3* Cat is injured directly by the explosion, whereas in *Quarry 4* Cat is injured by a landslide produced by a series of unusual events that are triggered by the explosion, not unlike the causal chain in *Palsgraf*. The intuitive difference between the cases remains even if we stipulate that Cat's ex ante probability of harm is the same in both. Perhaps the explanation as to why we should distinguish these cases is that, in *Quarry 4*, the relevant causal requirements are not satisfied because the causal route is too attenuated. Exactly what the relevant causal requirements are is of course a complex question that we cannot hope to settle here, but it should encourage us to revise *Nonrelational Fault* into the following:

*Nonrelational Fault 2*: When A commits some act,  $\phi$ , it constitutes a directed wrong against B if (1)  $\phi$  is either (i) at fault in relation to a third party or (ii) at fault tout court, (2)  $\phi$  causes harm to B, and (3) any relevant causal conditions are satisfied.

If highly unusual causal routes are ruled out, *Nonrelational Fault 2* can distinguish between *Quarry 3* and *Quarry 4*. Reflection on these cases shows that we should tread

carefully around *Palsgraf*. It is easy to confuse causation with noncausal elements of wrongdoing such as foreseeability, because unforeseeable outcomes are often those that arise via unusual causal routes. Nevertheless, causation and foreseeability are distinct elements of wrongdoing that should not be conflated. The important point for our purposes is that some of the concerns about broadness might be addressed once we factor in causal requirements. *Nonrelational Fault 2* can therefore accommodate the decision in *Palsgraf* if it can be explained in terms of lack of appropriate causation, although admittedly it is difficult to know whether this is so until we have a fleshed-out account of causation and a definitive analysis of *Palsgraf*, neither of which is likely to be forthcoming.

In any case, relying on causal principles will not solve the broadness problem completely. In fact, this problem is deeper than we have so far acknowledged. Once we accept that one can be wronged in the directed sense by an act that is at fault in relation to another, or to no one in particular, it is not obvious why we should not apply this view in cases where the reasons why the act is wrong have nothing to do with the outcome that occurs. Consider:

*Fraud*: Ali defrauds Bea and steals her money, leaving her destitute. He preys on people with no family whom he considers easier targets. Unforeseeably, Bea does have a relative, Cat, who is traumatized by Bea's experience and suffers from PTSD.

*Priceless Artifact*: Ali wants to blow up a priceless cultural artifact for fun. Doing so will not foreseeably harm anyone because Ali erects a fence. Cat is on a hiking trail but is led astray by paths created by animals, and through a gap in the fence that the animals have also torn down. Cat is injured in the explosion.

In *Fraud* and *Priceless Artifact*, Cat is an unforeseeable victim who is harmed by a wrongful act. In both cases, though, the grounds on which the act is wrong not only have nothing to do with the act's relation to Cat, but they also have nothing to do with the type of harm that Cat suffers. In *Fraud*, the wrong is grounded in the intentional violation of property rights rather than the infliction of psychological harm, and in *Priceless Artifact* Ali's act is not a directed wrong at all but a violation of impersonal values.

There are a few options here. The most expansive, and perhaps most radical, is to bite the bullet and insist that directed wrongs occur in all these cases. Though some may dismiss this out of hand, I think it is more arguable than it first appears. It is supported by the fact that, in all of them, the essence of the complaint articulated above still applies. In *Fraud* and *Priceless Artifact*, it is understandable that Cat would be angry at Ali, even though what makes the act faulty had nothing to do with the prospect of harm to Cat. It is still true that Ali was not morally free to harm Cat and did so anyway. It would be appropriate for Ali to experience guilt in relation to Cat rather than mere agent-regret. If he were to view himself as equivalent to someone who innocently caused harm, I suspect we would find something lacking in his response, a failure to take responsibility for the outcome of his actions.<sup>15</sup> The

<sup>15</sup>David Enoch has distinguished between being responsible and taking responsibility and suggested that we not only have the power to take responsibility for certain outcomes of our actions, but we are sometimes under a duty to exercise this power. If it is doubted that Ali is already responsible for the harm to Cat, cases like *Quarry 3* seem to me excellent candidates for Enoch's suggestion. Perhaps even more so than cases of

insufficiency of agent-regret does not depend on causing the kind of outcome that made the act wrongful; it depends on the fact that the agent is culpable for the wrong that leads to harm.

Bear in mind that I am not necessarily arguing that *Palsgraf* was wrongly decided, for two reasons. First, as we previously clarified, we are interested in directed wrongdoing in the moral sense rather than legal liability. There are many reasons why liability to compensate should be more restrictive in scope than directed moral wrongdoing; for example, to prevent it from becoming over-burdensome, disproportionate, or difficult to predict (we will return to these issues later).

Second, the decision in *Palsgraf* was correct as a decision about *the wrong of negligence*. The defendant in this case did not commit *this* wrong against the plaintiff. Instead, I am suggesting that they committed the wrong of harming through wrongful action, which has a distinct scope and grounds. Recognizing this wrong does not require us to reject the ruling in *Palsgraf* as a decision about negligence. The court might have been right not to enforce this other wrong, or to muddle it with the wrong of negligence. None of this means that *Nonrelational Fault 2* is false.

## VI. Sufficiently Similar Grounds

Biting the bullet will not appeal to everyone who accepts that Ali wrongs Cat in *Quarry 3*, so let us explore a narrower way of understanding the wrong in this case. We might say that Ali wrongs Cat not just because Cat is harmed by Ali's faulty act, but also because what happens to Cat is the same kind of thing that could have happened to hikers like Bea, which made Ali's act faulty in relation to them. The impact on Cat is sufficiently similar to the grounds on which Ali's act is faulty in relation to hikers that we are willing to accept that Ali's fault toward the hikers forms part of a wrong to Cat. By contrast, we have no complaint about being unforeseeably harmed by faulty acts if the reasons why those acts are faulty have nothing to do with their tendency to cause the outcome that in fact occurs. This is why we are uncomfortable with treating fraud or cultural vandalism as an element of the wrong of inflicting physical or psychological harm, while we are more willing to treat risking harm (or doing harm) to third parties as an element of such a wrong, as in *Quarry 3*. This amendment gets plausible results across the range of cases we have considered. It implies that Cat is not wronged in *Fraud* and *Priceless Artifact* but is wronged in *Quarry 3*.

We might then make the following variation to *Nonrelational Fault 2*:

*Nonrelational Fault 3*: When A commits some act,  $\phi$ , it constitutes a directed wrong against B if (1)  $\phi$  is either (i) at fault in relation to a third party or (ii) at fault tout court, (2)  $\phi$  causes harm to B, (3) the outcome for B is sufficiently similar to the grounds on which  $\phi$  is faulty and (4) any relevant causal conditions are satisfied.

---

innocent harm since, in the present case, the agent is already culpable for some form of wrongdoing. All the more reason, we might think, for them to take responsibility for the unforeseeable harms. See David Enoch, *Being Responsible, Taking Responsibility, and Penumbral Agency* in LUCK, VALUE, AND COMMITMENT: THEMES FROM THE ETHICS OF BERNARD WILLIAMS (Heuer and Lang eds., 2012) and *Tort Liability and Taking Responsibility* in PHILOSOPHICAL FOUNDATIONS OF THE LAW OF TORTS (Oberdieck ed., 2014).

*Nonrelational Fault 3* also distinguishes between *Quarry 3* and *Quarry 4*. In *Quarry 3*, Ali's act is faulty in relation to hikers like Bea, but nevertheless Cat is wronged even though it is unforeseeable that Cat would end up being harmed, because the harm to Cat (physical injury by explosives) is sufficiently similar to what grounds the faultiness of Ali's act (the threat of physical injury by explosives to hikers like Bea). In *Quarry 4*, Ali's act is also faulty in relation to hikers like Bea. However, Cat is not wronged because the harm Cat suffers (being hit by a landslide) is not sufficiently similar to what grounds the faultiness of Ali's act. In both cases Ali's act is wrong in virtue of the risk of explosive damage and not landslides, so he wrongs Cat in *Quarry 3* and not *Quarry 4*. Similarly, in *Palsgraf*, the defendant's act was wrong in virtue of the risks associated with pushing a passenger onto a moving train, not the risks associated with accidentally exploding fireworks.

One of the main challenges faced by *Nonrelational Fault 3* is to explain what counts as "sufficiently similar grounds" in a nonarbitrary way. This could mean that the *type* of harm suffered is the same although the specific way it comes about is different. This explains cases like *Fraud* and *Priceless Artifact*. In these cases, the outcome that occurs (physical or psychological injury) is of a different type to the potential harms that ground the initial wrong (deprivation of property and cultural vandalism). This may not be workable, though, because in *Quarry 3* and *Quarry 4*, the harm is of a similar type: injury to the body. To distinguish these cases, we would have to fall back on the idea that the relevant causal requirements are not met in *Quarry 4*.

Another possibility is that there are "sufficiently similar grounds" in *Quarry 3* because the way the harm occurs to Cat is as foreseeable as the way it threatens to occur to Bea. This casts the grounds on which an act is wrong more broadly, roughly as foreseeable setback to a significant interest. In *Quarry 3* Cat is harmed in the same way that Bea threatens to be harmed: physical injury through explosion. In both cases the route via which harm arises is foreseeable. Thus, Cat is wronged even though Ali's act is faulty because of the risk to Bea. This interpretation has the advantage of allowing some leeway in the precise causal route to the harm. Suppose that when Cat gets into the quarry, Cat does not end up close enough to be harmed by the explosion directly, but the explosion dislodges a rock that falls on Cat. If we think Cat is wronged by Ali in the initial case of *Quarry 3*, we should reach the same conclusion here because although the specific route to Cat's injury is different, it is just as foreseeable that they would be injured by debris as by the explosion itself. These variations of *Quarry 3* differ from *Quarry 4* where the route to the harm is unforeseeable since it involves a series of freak occurrences.

It is not clear that this is workable either, though, because the requirement that the specific route to harm be foreseeable is too restrictive. In the case of *Hughes v Lord Advocate*,<sup>16</sup> for example, a defendant was found liable for burns suffered by boys who had meddled with some equipment left out by Post Office workers. One of the boys tripped over a lamp which fell into a hole and exploded, causing him to fall in and burn himself. The House of Lords found in favor of the boys, Lord Morris holding that burn damage was foreseeable and the defendant should not escape liability just because they could not foresee the precise way in which it came about. This seems a sensible conclusion. If a person creates an obvious danger, they should not escape responsibility because the specific route to the harm was unforeseeable, as long as the

<sup>16</sup>[1963] AC 837.

outcome was foreseeable in a more general sense. A plausible gloss on the “sufficiently similar grounds” condition should therefore allow that the specific route need not be foreseeable while also holding onto the idea that some unforeseeable harm, such as that in *Quarry 4*, is ruled out.

The precise delineation of *Nonrelational Fault 3* will be complicated and no doubt there will be edge cases. However, these issues should not cause us to lose sight of the plausibility of *Nonrelational Fault 3*. Not only does it give us the right results across the range of cases we have considered, enabling us to distinguish between *Quarry 3* and *Palsgraf*, but it also runs parallel to principles frequently invoked in causation cases. For example, consider:

*Unusual Murder*: Ali wants to kill Bea, so Ali points a gun at Bea and pulls the trigger. Owing to an unforeseeable defect, the gun misfires and the bullet hits a wall, but it ricochets and then kills Bea.

Intuitively, Ali murders Bea despite the fact that the killing happens via an unforeseeable causal route. This is because Ali achieves his purpose in killing Bea in roughly the manner he intends, so the causal route that occurs is sufficiently similar to the predicted one. Compare this with:

*Lightning*: Ali shoots at Bea and misses, but Bea jumps away in fright and is struck by lightning and killed in the spot where she lands.

We can stipulate that the *ex ante* probability of the gun having the defect in *Unusual Murder* and Bea being struck by lightning in *Lightning* is the same. It is not the probability that distinguishes these cases but the fact that the causal route that occurs in the first case is sufficiently similar to the predicted one, and in the second it is not. Similarly, in *Quarry 3*, Ali’s act is negligent because it poses a risk of explosive damage. Although it ends up happening to an unforeseeable victim (Cat), the outcome is sufficiently similar to the risks that made the act wrong in the first place for us to conclude that Ali wrongs Cat. In this way, the “sufficiently similar grounds” condition can be seen as mutually supportive of the cause/coincidence distinction.

To sum up, then, we have seen three main arguments in relation to *Quarry 3*. First, it is appropriate for Ali to feel guilt toward Cat and inappropriate to feel mere agent-regret. Similarly, he has more reason to apologize for what he does to Cat than for an innocently caused equivalent harm. This argument supports *Nonrelational Fault 2*, the broad view that we always wrong those we harm via our wrongful actions, as long as relevant causal conditions are satisfied. The first argument also supports *Relational Fault 3*, the narrower view that *Relational Fault 2* must be satisfied in addition to the requirement that the outcome that occurs is sufficiently similar to the grounds on which the act is at fault. Two further arguments also support *Relational Fault 3*. One is that it gives us intuitive verdicts across the range of cases we have considered. It allows us to explain our initial verdict in *Quarry 3* without implying that directed wrongs occur in *Quarry 4*, *Fraud*, or *Cultural Vandalism*. Finally, the “sufficiently similar grounds” condition, though difficult to gloss precisely, is consonant with distinctions we commonly draw elsewhere in our thinking about wrongdoing, such as that between causes and coincidences. Of course, even those who feel the force of some or all these arguments might think the objections are still too weighty to accept them. In the next section I attempt to alleviate the weight of some of these objections.

## VII. Burdensomeness and Ex Ante Guidance

So far, I have been stating the argument in terms of directed wrongdoing. In *Quarry 3*, Ali wrongs Cat. This seems to me a natural way of stating the normative relation between Ali and Cat. There are other things we could say about this relation, for example (1) that Ali breaches a directed duty to Cat, (2) that Ali violates a right Cat holds against Ali, and (3) that Cat can demand some response from Ali such as apology or compensation. Some, such as Nicholas Cornell, argue that wrongs and rights violations come apart, so that it is possible to wrong someone without violating their rights. Accepting Cornell's view would be convenient for my purposes because in *Quarry 3* it is prima facie more plausible to say Cat is wronged than that Cat's rights have been violated. Accepting the latter claim might be a step too far for some.

Although Cornell's view is helpful for my purposes, I do not think it is right, at least in one respect. While I have no preference about how best to use the language of rights, Cornell identifies one normative feature of rights that, in his view, does not apply to wrongs: rights constrain ex ante conduct whereas the prospect of wronging someone needs not. To motivate this, Cornell presents several cases. One involves two children who are threatened with harm because A negligently fails to warn them about a fireworks display, and A can now warn only one of the children about the impending explosion. One child has a family who will be grief-stricken over the child's death; the other child's death will not cause grief to anyone. Cornell claims that we should not take these third-party harms into account when deciding whom to warn—they have no purchase on our ex ante reasoning. This shows that, though the parents of the child may be wronged, they lack rights against being wronged.

There are some issues with this example. The case should hold all else equal between the two children (there is no additional reason to save one or the other) and focus on the fact that only one death will result in grief. But by virtue of this fact, all else is not equal, since one child is living life with a family and the other is not. We would expect this to make the child without a family worse off, and it is plausible that preexisting well-being is relevant to decisions about whom to save. Another confounding feature is that both children will be wronged as the threat was created by negligence. Although one wrong may cause more harm, it does not follow that one child is wronged more than another. As a general matter, the severity of wronging does not track the severity of harm. We see this when we reflect that, though the murder of a child and the murder of a hundred-year-old cause different amounts of harm (because the child is deprived of more future life), they are equally wrong.<sup>17</sup> So it may be that our intuitions about Cornell's case are affected by the fact that the children will be equally wronged rather than the fact that third-party harms are irrelevant in calculations about whom to save.

To eliminate these concerns, consider the following:

*Third Party Harm:* Through no one's fault a trolley is on the loose that will kill either Cat or Bea. Ali can save either Cat or Bea at no cost but not both. All is equal between Bea and Cat except that no one will witness Bea being killed, whereas five people will witness Cat being killed and, as a result, suffer

<sup>17</sup>See Jeff McMahan, *The Ethics of Killing: Problems at the Margins of Life* (Oxford, 2002) 235 for the argument that killings are equally wrong even if the victims suffer different amounts of harm by dying.



debilitating PTSD for the rest of their lives. These outcomes are entirely foreseeable to Ali.

This case treats Bea and Cat as equal in all respects but the one we wish to focus on, namely the presence of third-party harms. Although the threat is no one's fault, Bea and Cat have a right to be rescued if this can be done at little or no cost. Some may be skeptical of a right to rescue, but here I mean only that there is a constraint on Ali's ex ante conduct. The fact that someone will die unless he performs a costless rescue makes him morally unfree to do anything else. If that is the mark of rights, Bea and Cat have a right to be rescued.

In *Third-Party Harm*, though Bea and Cat have equal rights to be saved, Ali ought to save Cat rather than Bea as this will prevent further serious harm to five other people. The debilitating PTSD of these five gives Ali some reason to save Cat rather than Bea, and since all else is equal, this tips the balance in favor of saving Cat. To strengthen this conclusion, simply increase the amount of third-party harm. Suppose that ten people will suffer PTSD, then thirty, a hundred, a thousand, and so on. At some point, it will become difficult to insist that the prospect of third-party harm is irrelevant for ex ante decision-making. This suggests not only that the third parties are wronged if Cat is not saved, but that their rights are violated too.

*Third-Party Harm* eliminates wrongdoing, but similar conclusions follow when we return to the structure of Cornell's original case. Consider a variation of *Third Party Harm* which is the same except that the trolley is loose due to Ali's negligence. Imagine again that five people will suffer PTSD unless Ali saves Cat, then imagine ten people, twenty, thirty, a hundred, a thousand, and so on. None of this affects the gravity of the wrong perpetrated on Cat or Bea—in both cases, they will suffer the wrong of being negligently injured or killed. But the third-party harms give Ali some reason to save Cat rather than Bea, and thus Ali violates their rights if he saves Bea instead.

It might be thought that in saying all of this I am arguing myself into a corner. After all, if wrongs against unforeseeable victims in cases like *Quarry 3* constrain ex ante conduct and are therefore also rights violations, does not this make the broadness problem even more acute? At least in Cornell's fireworks case, the third-party harms are foreseeable. When they are unforeseeable, as in *Quarry 3*, surely positing rights against their suffering is even more unsupportable. If unforeseeable victims of my wrongdoing have rights against me, this means the class of people who possess such rights is implausibly large and is, by definition, difficult or impossible to predict. The problem is most severe for *Nonrelational Fault 2*, but it also applies to *Nonrelational Fault 3*. On the latter view, the "sufficiently similar grounds" condition mitigates the broadness problem, but there is still a potentially large class of persons who might unforeseeably be harmed.

It is important to remember we are working with a specific conception of rights, one that correlates to duties that constrain ex ante conduct. Possessing a right of this kind does not imply a right in any other, thicker sense. Nevertheless, we might think this exacerbates the broadness problem for a few reasons. For one thing, if unforeseeable victims in cases like *Quarry 3* have rights, surely this is too many people for me to have in contemplation when I deliberate about how to act. Possessing a right means I must, in some sense, give due consideration to the right-holder, but there is a limit to the number of people who can feature in this way during my deliberation, otherwise, my practical reasoning would be hopelessly cumbersome. What is more, it is not even clear how I could take these potential victims into consideration. Since it is



unforeseeable that they might be harmed, there are no clear precautions that can be taken. In this respect, they are unlike the foreseeable victims of my actions where the foreseeability of harm gives me a guide as to what preventive measures I might take.

However, I do not think any of this is a problem for *Nonrelational Fault 2* or 3. To explain why, consider:

*Button:* If I press a button a random person in the world will be killed. I have no reason to press the button. I press the button.

Clearly, I act wrongly by pressing the button. To whom do I owe the duty not to press it, and who possesses a right that I do not press it? Suppose we say: Everyone in the world. This does not mean that the right is possessed by some collective entity. Rather, it is an aggregation of billions of individual rights. A right/duty pairing exists between me and everyone in the world despite the fact that the one in seven billion chance of dying is far less than the chance of third parties suffering harm in all of the cases we have discussed so far.

Does this answer make rights and duties overly burdensome or incapable of ex ante guidance? I suggest not, because recognizing them does not necessitate my having everyone in the world in contemplation, at least in any granular sense, when I decide how to act. Seven billion people possessing the right rather than one makes no difference at all to my ex ante practical reasoning. In both cases the content of my duty is the same: I must not press the button. There is no additional action that must be taken, or that I must consider at the practical deliberation stage, by virtue of a right possessed by so many people. All these things are also true in *Quarry 3*. In order to respect the rights of people like Cat who will unforeseeably be injured by the explosion, Ali need not have them in contemplation when deliberating. Moreover, because the duty to Cat piggy-backs on the duty to Bea and others like her, it is not more burdensome and does not require further action. Ali can respect Cat's rights by respecting Bea's. This is true of all rights conferred by *Nonrelational Faults 2* and 3. They can all be respected by respecting other rights that already constrain the agent's conduct.

There is another sense in which positing such rights might be too burdensome, which is that if people like Cat can claim compensation from Ali in *Quarry 3*, Ali's potential reparative burden is much greater than it would be if they were unable to claim. However, we must distinguish the question of whether a right has been violated in the sense that a duty that constrains ex ante conduct was breached and the question of what response is owed as a result. This version of the burdensomeness objection applies only to liability to compensate, but nothing I have said so far entails that those who commit directed wrongs against unforeseeable victims are always required to compensate.

Whether there is a duty to compensate depends on the circumstances. This is because compensatory duties—like all duties—are subject to a stringency threshold. There is a maximum burden a wrongdoer can be required to bear for the sake of fulfilling their duties. For example, if forcing them to compensate would make them destitute and would not benefit the victim a commensurate amount, the duty to compensate does not arise.<sup>18</sup> In cases where there are many victims, the stringency threshold is likely to be met more quickly. On the other hand, if Cat is Ali's only

<sup>18</sup>For defense of this claim, see Todd Karhu, *Proportionality in Liability to Compensate* 42 LAW & PHIL. 583 (2022) and Adam Slavny, *Wrongs, Harms, and Compensation*, 46.

victim in *Quarry 3* and compensating them would not be too costly for Ali, it is plausible that he has a duty to do this. So, the burdensomeness of liability to compensate is not a reason to reject the presence of directed wrongdoing and the violation of a right/duty pairing this implies. We can see this clearly by comparing compensation with less burdensome responses such as apology. In *Quarry 3*, if Cat is paragliding with twenty friends, all of whom are injured in the explosion, it is not implausible to think Ali has a duty to apologize to all of them, even if we do not think he ought to compensate all of them at overwhelming cost to himself. We can explain the lack of compensatory duties in terms of the stringency of Ali's corrective duties, while the duty to apologize remains because apologizing is less burdensome.

There are other reasons to limit compensatory duties, especially when it comes to legal liability. Liability rules must be applied across a broad range of cases in nonideal epistemic conditions. In particular, the obvious difficulty of predicting liability to unforeseeable victims can create problems for agents in effectively managing their potential liabilities, and when these problems are systemic, rule of law values can become eroded. We cannot appeal to the implausibility of over-broad legal liability as a reason to doubt *Nonrelational Fault 2* and *3* since there are plenty of other reasons to limit the liability of those who harm unforeseeable victims.

### VIII. Effacing and Double Counting

The cases we have been discussing bear a resemblance to transferred malice cases in criminal law. A standard transferred malice case is like this:

*Transferred Malice:* Ali intends to kill Bea. He shoots in her direction but misses and kills Cat instead.

According to the doctrine of transferred malice, Ali is guilty of murdering Cat even though he did not intend to kill them. His intention to kill Bea is “transferred” onto Cat. This doctrine is similar to *Nonrelational Faults 2* and *3*.

The doctrine of transferred malice is subject to some objections, and it is worth considering whether they also apply to *Nonrelational Faults 2* and *3*. One problem with transferred malice outlined by Patrick Tomlin is that, if Ali murders Cat, the doctrine mislabels the wrong involved in a way that erases Bea from the picture.<sup>19</sup> If Ali's intention to kill Bea is “used” as an element in the wrong against Cat, making it the single wrong of murder rather than the dual wrongs of manslaughter and attempted murder, then the wrong against Bea is erased from the picture. Insofar as we are interested in the accurate labeling and recognition of wrongdoing, this is a problem. This objection might be raised in *Quarry 3* and perhaps gives us further reason to doubt *Nonrelational Faults 2* and *3*.

This objection does not apply to *Quarry 3*, though, due to one important difference. According to our analysis, there are two wrongs, one to Bea and one to Cat. Bea is not erased from the picture because the wrong to her is not “used up” in any sense—it exists in addition to the wrong to Cat. Problem solved? Not quite. Now we run into another worry: Double counting.<sup>20</sup> If *Nonrelational Fault 2* or *3* is correct,

<sup>19</sup>Patrick Tomlin, *Accidentally Killing on Purpose: Transferred Malice and Missing Victims* 41 LAW & PHIL. 329 (2022).

<sup>20</sup>See Findlay Stark, *Recycled Malice* 82(3) *The Cambridge Law Journal* 509 (2023).

the wrongful risk imposed on Bea in *Quarry 3* is counted as part of the wrong committed against Cat. But if we are also to say that Ali wrongs Bea, this wrongful risk is counted twice, once for the wrong to Bea and once for the wrong to Cat. Perhaps it is unfair that Ali's wrongful risk should be counted twice in this way.

The double counting objection is persuasive in contexts where we want to maintain a tight proportional relationship between the severity of the wrong and the burden the wrongdoer faces as a result. This is true when issuing punishment.

Being punished twice for the same malicious intent risks heaping a disproportionate burden on the wrongdoer. Although I previously noted that all duties have a stringency threshold, so it is possible for compensatory duties to become disproportionate, it remains the case that the burden of repair need not be tightly proportional to the severity of the wrong in the same way that the burden of punishment should be.

Suppose there are twenty people in the quarry when Ali sets off the explosion. In this case, Ali would wrong each of them and they would all have a claim for compensation against him. This is true despite the fact that Ali's negligent act might have been no more culpable or unreasonable than if just one person had been injured. In a sense, his unjustified act is "counted" multiple times for each claim that can be made against him, but this is no bar to recognizing all these claims. Unless the reparative burden is so large that it exceeds the stringency threshold of Ali's duties, there is no objection. In other words, the double counting objection applies to certain responses to wrongdoing such as punishment. It does not apply to other responses such as apology and compensation, or to the prior claim that a directed wrong has been committed.

## IX. Conclusion

In this paper, I have presented two ways of understanding the wrong done to Cat in *Quarry 3*: That Cat is harmed by a wrongful action, and that the outcome for Cat is sufficiently similar to the grounds on which the act that harmed Cat is at fault. Accepting either of these ideas requires significantly revising the *Relational Fault Requirement*, but they nevertheless maintain the kind of moral relation between agent and victim necessary for directional wrongdoing. They are both principles that shape our interpersonal moral relations, and they can vindicate the sense in which the wrong to Cat in cases like *Quarry 3* is directional without yielding to Lewinsohn's conventionalist challenge.

**Acknowledgements.** I would like to thank the participants of the Centre for Ethics, Law, and Public Affairs at Warwick for their valuable comments. Special thanks go to Adam Perry, David Enoch, Kate Greasley, James Edwards, and Sandy Steel for organizing the OxJuris workshop, which yielded valuable feedback on the paper. Thanks to all the participants of this workshop, and especially to Alex Kaiserman for his commentary and Ruth Chang for her email correspondence.

**Competing interest.** The author declares none.