

## VIII—SITUATIONAL DEPENDENCE AND BLAME'S ARROW

JESSICA ISSEROW

A looming deadline. A difficult situation at home. A heated phone conversation that redirects our attention. Certain features of our circumstances can be (at least partially) excusing; sometimes, agents who act wrongly in the face of circumstantial pressures are not (that) blameworthy for having done so. But we're rather bad at detecting these factors that excuse others from blame. When put together, these two observations yield an under-appreciated problem: we fall short of procedural norms of blame in fairly systematic ways.

### I

*Introduction.* Two observations about our moral lives generate an important problem when put together:

*Excusing Circumstances.* Certain features of people's circumstances can be (at least partially) excusing; even if they act wrongly, they're not (that) blameworthy for having done so.

*Blame's Blind Spot.* We're unreliable at detecting features of others' circumstances that (at least partially) excuse them from blame.

The second observation is of course an empirical one, but it's well-supported. The first reflects a philosophical claim, and I shall argue that the case for this is similarly strong. When put together, these observations yield an under-appreciated problem: we fall short of procedural norms of blame in fairly systematic ways.<sup>1</sup> The procedural norms in question are the following ones:

---

<sup>1</sup> I'm not the first to notice that people's vulnerability to circumstantial pressures carries implications for the ethics of blame. However, those who have so far teased out these implications have focused primarily on *standing* norms rather than *procedural* ones (see Isserow 2022; Piovarchy 2023).

*Epistemic Norm.* A blamer's belief that a blamee is blameworthy should clear some epistemic bar (for instance, justification or knowledge).

*Proportionality Norm.* A blamer's blame should be proportionate to a blamee's blameworthiness.

I'll set out by defending the above observations (§§II–III). I'll then elaborate upon precisely what procedural norms of blame are, and why we should take them to be important (§IV). Finally, I'll expose the problems that emerge once we put our two observations together (§§V–VI).

## II

*Excusing Circumstances.* My first task will be to defend Excusing Circumstances. I'll begin with the strong intuitive case in its favour before diving into any of the underlying philosophy in earnest. Consider the following pairwise comparisons:

*Dog-Walker 1.* I'm walking my dog Perdita on a Saturday afternoon. She does her business on your front lawn. I clean up the dog-droppings.

*Dog-Walker 2.* I'm walking my dog Perdita on a Saturday afternoon. I receive a phone call from my draconian boss Cruella. She insists that I come into the office, and threatens to fire me if I don't. (She doesn't truly mean it; I know this to simply be her way of communicating that she will otherwise be very cross.) Cruella's call inspires a mild panic attack, and I quickly pull Perdita along. Upon doing so, however, I discover that Perdita has done her business on your lawn. Given my state, I can't bring myself to waste any time. I abandon the dog-droppings.

*Jogger 1.* I'm out running, and see an old man struggling to pick up his wallet. I stop to help.

*Jogger 2.* I'm out running after having just received news about a sick relative. My mind is buzzing, making plans. (How soon can I get there? Can I reschedule work?) My eyes quickly pass over an old man reaching for something on the ground. I'm distracted, and don't reflect on the scene any longer. I don't help.

*Promotion 1.* Jack is out for drinks with colleagues to celebrate his promotion. One colleague, John, remarks that Jack only got the promotion because the boss thinks that Jack is sexy. Jack is a little put off by the remark, but he laughs it off.

*Promotion 2.* Jill is out for drinks with colleagues to celebrate her promotion. One colleague, John, remarks that Jill only got the promotion because the boss thinks that Jill is sexy. Jill's colleagues often make remarks like this, and John does so with irritating frequency. Jill is fed up. She takes off her shoe and hurls it at John's head.

Given their circumstances, the second agent in each pair encounters obstacles to acting morally well that the first doesn't. I am, for instance, emotionally distressed in *Dog-Walker 2* and distracted in *Jogger 2*. And given the wider context, Jill is provoked by John's remark in a way that Jack isn't. These obstacles seem to go some way towards explaining why we'd be reluctant to hold these moral mishaps against the second agent in each pair, or to judge them too harshly on account of them—there's the inclination to say that each has an *excuse*.

These efforts at intuition-mongering are of course only a start. It will be helpful to explain how exactly I am understanding excuses in this context, and why it's plausible to suppose that these agents do indeed have them. We can start by distinguishing different questions that we might want to ask about *Dog-Walker 2*. First, have I *acted wrongly* in abandoning the dog-droppings? The answer here is plausibly yes. But there's a second question that concerns the agent more than it does the act: *am I morally responsible* for this wrongdoing?

The answer to this second question is less clear. Given the stress inspired by Cruella's phone call, we may want to deny (or challenge the extent to which) I'm responsible for this wrong. If my responsibility was indeed diminished, then I have an *excuse*. Excuses, then, prevent a particular sort of judgement about the moral status of our actions from being paired with a particular sort of judgement about our agency; they sever (or at least weaken) the connection between wrongdoing and responsibility (Austin 1957, p. 2; Brink 2013, p. 131; compare Sliwa 2019, pp. 63–6).

Much of what I have to say about excuses will be couched in terms of blameworthiness rather than responsibility, and I will be helping myself to a cluster of assumptions about both. While it's worth

making these assumptions explicit for clarity's sake, I don't take the case for Excusing Circumstances to stand or fall with them. Just about everyone should take something like Excusing Circumstances to be true, even if they disagree with me about precisely how we ought to account for it.

One assumption that I'll be making concerns blame itself. As I (and indeed, many others) see it, blame is crucially a matter of experiencing *reactive attitudes* of some kind; for instance, resentment, anger or indignation (Strawson 1962). I'll also be helping myself to the (somewhat simplifying) assumption that an agent is blameworthy if and only if she is responsible for wrongdoing. Importantly, responsibility for wrongdoing reflects the conditions under which blame would be *fitting or deserved*. It's a further question whether blame would be appropriate or justified, all things considered (see §IV). The order of explanation here should also be read from right to left; an agent is blameworthy for a wrong (when she is) *in virtue of* being responsible for it (see Brink and Nelkin 2013, pp. 287–8). The factors that undermine an agent's responsibility for wrongdoing, then, will by extension tend to undermine her blameworthiness as well. There's room to bicker with each of these claims, which have garnered supporters and detractors alike. But every paper needs to take some assumptions for granted, and I won't waste space here defending my own.

Now on to excuses. Four features help to set these apart from other normative phenomena. First, excuses attach only to wrongdoing. In this respect, they differ from *justifications*, which attach to actions that are morally right or justified, all things considered. Excuses don't, then, undermine our judgement that an agent acted wrongly. Indeed, the excuse presupposes that she did act wrongly—that there is something *to excuse*. Justifications, by contrast, do undermine our judgement that an agent acted wrongly; if her act was morally justified (on the whole), then it was *ipso facto* not wrong.<sup>2</sup>

Second, excuses can only be offered on behalf of agents who have the capacities required for moral responsibility, and who are therefore generally apt targets for praise and blame. In this respect, excuses differ from *exemptions*, which are offered on behalf of

---

<sup>2</sup> This way of drawing the justification–excuse boundary is common (Brink 2013, p. 131; Austin 1957, p. 2). But it is not the only game in town: see Wallace (1994) for an alternative paradigm.

agents who *lack* the capacities needed for moral responsibility, and who are therefore not apt targets for praise and blame. A tactless remark about the size of a dinner guest's bottom might be blameworthy coming from me. But it plausibly wouldn't be coming from a young child or from someone with advanced dementia.<sup>3</sup>

Third, excuses sometimes call an agent's *direct* responsibility for a wrong into question without undercutting her *indirect* (and thus *overall*) responsibility for it. Consider Linda, who applies lipstick while driving. Being distracted, she doesn't notice a child on the road, whom she subsequently hits and injures. Linda isn't *directly* responsible for hitting this child; she didn't know they were there, and she certainly didn't hit them intentionally. Her ignorance of the child's presence thus looks like an excuse. But Linda is not excused *tout court*; for we can *trace back* this ignorance to an earlier choice for which she *was* responsible—namely, choosing to apply lipstick while driving. Linda's ignorance, then, is not really an excuse; it doesn't shield her from indirect (and thus overall) responsibility for hitting the child. When I speak of excuses, I'll assume that we're not dealing with any such 'tracing cases'.

Finally, excuses are scalar—that is, they admit of degrees (Brink 2013, p. 134; compare Sliwa 2019, p. 69). One may have a partial excuse for wrongdoing, and thus be less blameworthy than one would otherwise be. Or one may have a (full) excuse for wrongdoing, and thus be fully excused from blame.

In so far as excuses temper moral responsibility (and, by extension, blameworthiness), it stands to reason that they will typically call into question the extent to which an agent satisfies the conditions for moral responsibility (and hence the extent to which she is blameworthy). On one prominent model, the conditions for moral responsibility are twofold: in order to be responsible for her actions, an agent must (i) have sufficient *control* over them, and (ii) in some sense (even if not the fully literal one) *know* what she is doing. These are, respectively, the control and epistemic conditions on moral responsibility.

<sup>3</sup> I won't take a stand on precisely where to draw the excuse–exemption boundary—that's a tricky question, and answering it won't be necessary for my purposes. For further discussion, see Brink and Nelkin (2013, p. 291), which effectively takes exemptions to be the limit cases of excuses.

With this background in place, it may be thought that it's going to be embarrassingly easy to demonstrate that the agents in our case studies have an excuse; we need only establish that they plausibly lack sufficient control over their actions, or that they fail to satisfy some responsibility-relevant epistemic standard. As far as the epistemic condition goes, I think this is largely true. While there's disagreement as to whether *moral* ignorance excuses, it's generally agreed that *non-moral* ignorance—roughly, unawareness of morally relevant descriptive facts—can. And the case study that invokes ignorance as an excuse (*Jogger 2*) trades on ignorance of this latter, non-moral sort. In so far as I fail to properly see the person in need, I'm not plausibly blameworthy for failing to help them. We might substantiate this diagnosis by appealing to the well-known phenomenon of 'looking but failing to see'. Drivers often fail to *notice* threats on the road while on their phone (and are hence more likely to have an accident) even when they *see* them (see [Piovarchy 2022](#), p. 2017). Being distracted may not ground full-scale ignorance, but it's arguably enough to call the epistemic condition into question.

Matters are more complicated as far as the control condition is concerned. 'Can' is notoriously slippery; it's easy to generate excuses willy-nilly by insisting that an agent simply *couldn't* have acted as she should have—and in *some sense* of 'can', these claims will, one expects, more often than not be true. Thankfully, the cases on which I rest my arguments aren't especially controversial, and so shouldn't raise any suspicions that I am here simply stretching the meaning of 'can'. Many would likely allow that control has been compromised in these cases.

Consider *Dog-Walker 2*. Given the mild panic attack that Cruella's phone call sets into motion, my control is plausibly compromised, even if it's not eliminated completely. I'm *far less* blameworthy for abandoning the dog-droppings here than I would have been had these circumstantial pressures been absent, and so it's plausible that I at least have a *partial* excuse for what I've done. We can say something similar about *Promotion 2*, where there is likewise a challenge arising from the volitional side of things. Jill simply snaps. She *loses* control. So it's plausible that she's at least partially excused as well.

Note that it's important to distinguish a loss of control from other phenomena in the psychological neighbourhood. Jill doesn't fall prey to *weakness of will*; she doesn't simply *give in* to the temptation to hurl her shoe at John's head. (That would be understandable,

perhaps, but it would be no excuse!) Jill has every intention of keeping herself together. But her rage makes it incredibly difficult to *act on* that intention. It's in virtue of her *capacity to act on* her intentions having been compromised that she has an excuse (see [Brink 2013](#), pp. 131–3; [Sliwa 2019](#), pp. 55–7).

None of this is to suggest that Jill has a *full* excuse. My claim here is not that acts born of rage are never blameworthy. But they are, I think, at least *less* blameworthy than they would be if carried out in a calm and measured state of mind. It also helps to recall that I am working with the assumption that our scenarios are not 'tracing cases'; it is not as though Jill has a history of violent outbursts and has failed to bring her anger under control. If that were so, then we likely *would* judge Jill (indirectly) blameworthy for her shoe-hurling.

### III

*Blame's Blind Spot.* My second task will be to defend the following empirical claim:

*Blame's Blind Spot.* We're unreliable at detecting features of others' circumstances that (at least partially) excuse them from blame.

To this end, I'm going to draw upon a well-documented phenomenon known as the Actor–Observer Asymmetry in Attribution (AOAA):

We (actors) favour situational explanations for our own behaviour and we (observers) favour personal explanations for others' behaviour.

Psychologists generally adhere to more or less the same patterns when distinguishing situational factors from personal ones. Among the former are factors such as 'chance, the weather, a stimulus, another person with whom the actor interacts'; the latter include 'effort, ability, attitudes, personality, mood, desires' ([Malle 2006](#), p. 896). Some allude to the boundary of 'the human skin' as a heuristic: 'On sunny side of the epidermis are the ... situational forces that press inward on the person, and on the meaty side are the internal or personal forces that exert pressure outward' ([Gilbert and Malone 1995](#), p. 21). While this doesn't quite approach the level of precision that will tend to please philosophers, it will be enough for us to go on.

The observer side of AOAA will be more important for my purposes. To cite just some examples:

*Castro Case* (Jones and Harris 1967). Subjects are randomly assigned to deliver either a pro-Castro or anti-Castro speech. Even when listeners knew about the random nature of the assignment, they still attributed the views expressed to the speaker.

*Jury Case* (Kassin and Sukel 1997). Mock jurors infer guilt from confessions, even when they know that coercive means were used to elicit the confession.

Importantly, there's evidence for the *asymmetry* as well; we don't tend to discount situational factors when it's our own behaviour that's at issue:

*Driving Case* (Flick and Schweitzer 2021). When actors imagine themselves being involved in a car accident under both favourable and unfavourable driving conditions, they're more likely to blame the accident on situational factors when driving conditions are unfavourable. But when they (qua observers) imagine others being involved in a car accident, they're more likely to blame the accident on the driver, regardless of what driving conditions are like.

It will be helpful to introduce some qualifying remarks about AOAA before proceeding to draw any philosophical lessons from it. First, AOAA is not to be confused with the *Fundamental Attribution Error* (FAE) or 'Correspondence Bias', which is often said to reflect the observer side of the actor–observer asymmetry. FAE is (frustratingly) used in different ways. But it is, I think, often best thought of as picking out a particular instance of AOAA rather than as simply co-extensive with the observer side of it. For example, FAE often refers to a tendency for observers to *infer character traits* from behaviour. But AOAA need not involve any such inferences; it merely concerns the *kinds of explanations* offered for behaviour (see Malle 2006, p. 896). The mock jurors in Kassin and Sukel (1997) study might explain a coerced confession by appealing to the defendant's guilt (a 'personal factor') without inferring anything about the defendant's character traits.

Second, AOAA may well be explained by different underlying processes. One common explanation is that, in our capacity as



observers, we fail to properly *recognize* the situational factors that influence others' behaviour. But there are others (see Gilbert and Malone 1995). Even when we do recognize situational factors, our judgements are often skewed by unreasonable and unrealistic expectations about how a typical person would react in that situation. As Piovarchy (2022, pp. 2010–11) observes, background beliefs tend to skew judgements of capacity. Sometimes, it is so obvious to us (qua observers) which course of action is the right one that we find it difficult to imagine anyone failing to see this. Relatedly, and as Sliwa (2020) points out, there is evidence that we're susceptible to 'epistemic egocentrism' in our attributions of blameworthiness (Royzman et al. 2003): we often import our own privileged knowledge when judging what knowledge was available to others. This 'curse of knowledge' may lead us to attribute more insight to people than they actually had. So a failure to appreciate the operation of circumstantial factors is certainly not the only mechanism that underlies AOAA. But it will be the one that I focus upon in the interests of establishing my basic point. Narrowing my focus in this way should not affect the plausibility of my arguments.<sup>4</sup> If anything, it undersells just how widely the problem that I identify applies.

Finally, AOAA likely forms part of a wider self-serving bias. Watson's (1982) meta-analysis claimed strong support for AOAA, but Malle's (2006) more recent one has questioned the extent of that support. One notable issue is that AOAA doesn't seem to apply indiscriminately across all kinds of behaviour; the *valence* of the behaviour matters. Whereas we (qua actors) favour situational explanations for our own *negative* behaviour and (qua observers) favour personal explanations for others' *negative* behaviour (as AOAA predicts), this effect is, remarkably, reversed for *positive* behaviour: we (qua actors) favour personal explanations for our own positive behaviour, but not for the positive behaviour of others. Simply put: when *I* do something crummy, I'm a victim of circumstance, but when *you* do something crummy, it's because there's something crummy about you. However, when *I* do something fantastic, it's because *I'm* fantastic, while your fantastic feats don't necessarily say anything fantastic about you. Following Malle (2006, p. 896), this latter effect

<sup>4</sup> It does affect the accuracy my labelling, however. 'Blame's blind spot' suggests that it's *our failure to notice* situational influences that gives rise to the problem. And while that is my primary focus, there are in fact other underlying causes.

(which is better supported than the more indiscriminate version of AOAA) seems to reflect a kind of *self-serving bias in attribution*—a bias that arises from an interaction between the perspective taken (actor or observer) and how the behaviour is judged (positive or negative).

Here, at last, is the problem. We've seen that circumstantial factors can (at least partially) excuse bad behaviour. But AOAA suggests that we're rather bad at picking up on them. What this means, in effect, is that we're rather bad at picking up on factors that are crucial for deciding whether a given individual is blameworthy.

To demonstrate, consider ignorance cases like *Jogger 2*, where I 'look at but fail to see' someone in need, owing to circumstantial pressures that distract me. We should already expect blamers to have trouble detecting or fully appreciating ignorance (recall the 'curse of knowledge'). And AOAA should leave us feeling even more pessimistic in this regard; generally speaking, it seems that we cannot count upon blamers to give much consideration to the possibility that circumstantial factors may have induced ignorance and yielded an excuse. Or consider control cases like *Promotion 2*. Here, Jill snaps and loses her temper, owing to situational pressures that provoke her. AOAA should leave us feeling pessimistic that any onlookers will seriously consider the possibility that such pressures led Jill to lose control, thereby partially excusing her.

To put the point in a slogan, then: *blame's arrow doesn't fly true*. And this very much seems to be owing to procedural defects in how we go about blaming others. We're prone to blaming others without properly attending to a class of factors that are crucial for deciding whether (or to what extent) they *are* blameworthy.

Some may wonder why I'm restricting my focus here to *blame's* blind-spot.<sup>5</sup> The self-serving bias does after all work in both directions. In instances of *negative* behaviour, our under-appreciation of others' circumstances means that we neglect factors pertinent to assessing their *blameworthiness*. But in instances of positive behaviour, our under-appreciation of others' personal characteristics might mean that we also neglect factors pertinent to assessing their *praiseworthiness*. The comparative contribution of the person—as opposed to their circumstances—is generally thought to be important in deciding whether or not they are praiseworthy; inasmuch

<sup>5</sup> Thanks to Justin Snedegar and Berislav Marušić for pressing this point.

as someone is praiseworthy, their acting well cannot simply come down to luck (Sliwa 2016; Isserow 2019). In so far as we offload the explanatory burden for others' good behaviour onto their circumstances, then—as AOAA suggests we do—our judgements of praiseworthiness may tend to be off as well. This suggests that we may well be prone to failing to praise the praiseworthy as well as to blaming/over-blaming the non-/not-so-blameworthy—keenly generous when it comes to doling out blame, but rather miserly when it comes to dishing out praise.

I do suspect that there's a wider problem afoot here; perhaps our responsibility assessments more generally are being sung off-key. But I'll restrict my focus to blameworthiness in the interests of moral triage. It should, I think, worry us far more if blame's arrow doesn't fly true than if praise's doesn't. (The punishment of the innocent is of greater concern than a failure to reward the achieving.) This isn't to deny that there can be injustices in the distribution of praise as well (see Holroyd 2021), but I leave examining those for another occasion.

#### IV

*The Ethics of Blame.* Many inquiries into blame have tended to focus on the consumption side of things; the question of interest has typically been what makes an agent on blame's receiving end blameworthy. In recent years, however, more philosophers have shifted their focus to blame's production side; their question concerns the ethics of *blaming*. Even if someone *is* blameworthy, it doesn't necessarily follow that blame is called for—or that *our* blame is called for, anyway. Sometimes, it's simply not our place to blame. And even if it is, there are surely better and worse ways to do it.

There are two important theoretical divisions here with which we should familiarize ourselves. The first is between *standing* norms and *procedural* ones. The former caution against blaming when it isn't our place to do so. The 'non-hypocrisy norm' for instance, cautions against blaming others when we're guilty of parallel faults (Smith 2007, pp. 479–80; Todd 2019). *Procedural norms*, by contrast, advise against blaming *in particular sorts of ways*. The Epistemic Norm discourages being uncharitable in our blame (Coates and Tognazzini 2012, p. 204), whereas the Proportionality Norm discourages being too extreme (and perhaps also being too lax) with it (Smith 2007, p. 480).

Procedural norms are often motivated by the observation that blame has a punitive character: in blaming others, we place certain burdens and harms on them. Familiar procedural considerations that bear upon retributive justice are, then, often thought to bear upon blame as well. Coates and Tognazzini (2012, p. 204) note that epistemically sub-standard blame can reflect ‘a lack of adequate concern for moral innocence’, drawing parallels with the legal norm that a defendant’s guilt must be proven beyond reasonable doubt (2012, p. 204 n. 10). Likewise, the Proportionality Norm gets its moral heft from the idea that the punishment ought to be proportionate to the crime: minor transgressions don’t merit the same kind or degree of blame as more serious ones do.

Importantly, a violation of procedural norms can wrong others—even if they are indeed blameworthy and we have the standing to dole it out. If I am blameworthy, then I am *deserving of* blame. But that doesn’t yet decide whether it’s *permissible*, all things considered, to blame me—or to blame me to a certain degree. If my blameworthiness is in doubt (perhaps you’re unsure whether my volitional capacities were compromised), then you’d plausibly wrong me by blaming me. And if I’m only blameworthy to some small degree (suppose I forgot to reimburse you for lunch), then you’d likewise seem to go wrong by reacting as though I’d just been embezzling funds.<sup>6</sup>

A second important theoretical division is between the ethics of *private* and *expressed* blame. Procedural norms are plausibly relevant to both.<sup>7</sup> We would seem to wrong others in so far as we privately

<sup>6</sup> One need not take these procedural norms to be absolute ones. Better, I think, to interpret them as saying that there are strong *pro tanto* moral reasons against over-blaming, or against blaming in an epistemically sub-standard way. These reasons are, to be sure, capable of being outweighed; perhaps over-blaming a wrongdoer will be important for advancing social justice in certain contexts. However, it also needs to be acknowledged that the considerations that count against over-blaming or blaming on the basis of shoddy evidence are morally powerful ones—so they should not *too easily* be overridden either. Just as we would be hesitant to over-punish someone in the service of some cause, so too should we generally be wary of over-blaming others for the sake of some other good.

<sup>7</sup> Though perhaps they are relevant in different ways. Public blame has particular functions and features that are lacking in private blame; for instance, it plausibly serves as a kind of enforcement mechanism (detering bad behaviour), as well as a communicative one (advertising one’s moral commitments). And these may well make a difference to (for example) the kind of epistemic bar that needs to be cleared for public—as opposed to private—blame to be appropriate. It may turn out that the procedural norms on private blame are slightly different to those governing public blame. If that is so, then the assumptions with which I work here will inevitably be simplifying ones. Thanks to Hallvard Lillehammer for raising this point.

over-blame them, or privately blame them on the basis of shoddy evidence. Still, we shouldn't overlook the fact much private blame has no real practical upshots for its targets. Many of the harms associated with procedurally defective blame are owing to these hostile attitudes having made themselves known. So I shall focus primarily on expressed blame below.

## V

*The Epistemic Norm.* I want to propose that we fall short of:

*The Epistemic Norm* (EN). A blamer's belief that a blamee is blameworthy should clear some epistemic bar (for instance, justification or knowledge).

Our susceptibility to AOAA suggests that we're not as reliable as we should like when it comes to detecting circumstantial influences that can undercut attributions of blameworthiness. And this, in turn, suggests that we're falling short of EN. Situational factors can excuse, but we have a blind spot when it comes to picking up on them—there is, if you like, a glitch in our blameworthiness detection mechanism. We should therefore worry that our blame isn't living up to epistemic standards.

At this stage, one may feel entitled to ask: *which* epistemic standards? Surely we need to know what exactly the epistemic bar *is* before we can despair at having fallen short of it! But it's not clear that the foregoing argument is standard-dependent in quite this way. The essence of the problem is that we have positive reasons to doubt the reliability of a belief-forming process—namely, the process by which we arrive at beliefs about others' blameworthiness. And this problem seems to persist whether we take the epistemic standard to be knowledge or justification (on at least many understandings of these phenomena, even if not all).

Suppose the standard is *knowledge* (see [Littlejohn 2020](#); compare [Kelp 2020](#)). We might then worry that the relevant belief-forming processes are too unreliable to earn the imprimatur of knowledge. Or we might take AOAA to issue an undercutting defeater that precludes our blameworthiness-beliefs from qualifying as knowledge. Alternatively, suppose the standard is *justification* ([Rosen 2004](#)). We might then worry that our blameworthiness-beliefs are not properly

responsive to our evidence, or that we're failing to put together our evidence in the right way.

Of course, there's room to question just how much of a problem this really is. Perhaps we fall short of EN as far as strangers or acquaintances are concerned. We don't typically pause to consider whether the person who pushes in front of us at the bar might have been desperate to impress an employer, or whether the colleague who writes a snappy email response might have been exhausted after a night spent awake with a sick child. When it comes to our nearest and dearest, however, it's questionable whether we really do fall short of EN—for we have ample subject-specific information to draw upon here. As Sliwa observes:

[O]ur assessment of excuses becomes both more nuanced and idiosyncratic, the more familiar we are with someone and their quirks. We might not take offense at our partner's snide remark in the car because we know that she gets particularly stressed out driving in heavy traffic. We recognize this as an excuse for her misstep even though we would not recognize it as an excuse in general. (Sliwa 2019, p. 60)

However—and the high intuitive plausibility of this suggestion notwithstanding—it's possible that we overestimate ourselves in this regard. Research suggests that we are in fact prone to AOAA even when it is our intimates whose behaviour we are judging (Malle 2006, p. 904).

But let's simply grant that we're only usually prone to blaming strangers in an epistemically sub-standard fashion. Still, I'm inclined to ask: why is that *not* a problem? This push-back seems to presuppose that such blame oughtn't to be concerning. And that presupposition strikes me as questionable. If anything, we might be *more* worried about extending uncharitable blame to strangers. (Particularly so if those strangers are defendants in court.) Within personal relationships, the other person at least has the opportunity to *respond*. This is less likely to be true of strangers in passing interactions, or of acquaintances who may feel uncomfortable offering up personal information about themselves. Following Baron, excuses aren't always offered as declarations of non-responsibility, but often constitute invitations to temper blame—'pleas not to be judged harshly, or not to be misunderstood' (Baron 2007, p. 30).

Since strangers and acquaintances don't have similar opportunities to make these pleas, they seem more liable to being misjudged.

There's also room to question whether these arguments really *do* show that we're unreliable at forming beliefs about others' blameworthiness. Perhaps the problem is not that we are too harsh on others, but that we are too easy on *ourselves*.<sup>8</sup> I'm not averse to the suggestion that we might be too easy on ourselves; maybe we *are* too quick to identify situational factors as excusing when it is our own behaviour that we are judging. But the evidence does suggest to me that we are too harsh on others as well. If someone's pro-Castro speech was the result of a random allocation, then presumably it doesn't provide any new information about their feelings about Castro. But subjects don't properly attend to these situational factors, and so they assign more evidential weight to the speech than is warranted. Moreover, the evidential value of a confession is surely diminished if it was coerced; it is because subjects don't properly take into account the circumstances under which these confessions are made that they take them to provide more evidence of guilt than they should.

Finally, some might object that my arguments only have mileage because I've set the epistemic bar for blame too high.<sup>9</sup> Consider Dover's suggestion that reactive attitudes can serve as 'catalysts for further communication' (2019, p. 397). The resultant communicative exchange may be swift; blame may simply prompt its target to offer an excuse or an apology. Or it may continue longer still, providing opportunities for 'criticism conducted in a ... dynamic, egalitarian, back-and-forth mode' that yields 'substantive moral and interpersonal insight' (2019, p. 403). In keeping with this, we might want to view (expressed) blame as an opening to a conversation. And since that conversation is potentially valuable, we might want to *resist* claiming that an opening move can only be justified if our beliefs about others' blameworthiness are backed by sufficient evidence, or reflect knowledge. Perhaps we need only have some reason for believing the person blameworthy.

---

<sup>8</sup> Thanks to Will Gamester and Jessica Leech for raising this possibility.

<sup>9</sup> I thank previous audience members who have raised this style of challenge, including (if my notes and memory serve me right) Olle Blomberg, Jessica Brown, Sebastian Köhler, and Theo Murray.

I agree that others' wrongdoing can and often should prompt some kind of conversation. But I disagree that *blame* is always the best way to instigate such a conversation. Suppose I suspect on the basis of weak evidence that a friend betrayed my trust. One way to go about things would, of course, be to grow visibly irritated and accuse her. ('How dare you tell everyone my secret!') But surely a better, more morally respectful way to approach the matter would be to simply ask her: 'Did you happen to tell anyone about my secret?' We often (rightly) feel insulted when others jump to conclusions about our blameworthiness. And it's surely possible to create the space for critical discussion without leaping to such conclusions on the basis of little evidence. One expects that this will often make the ensuing discussion more constructive.

## VI

*The Proportionality Norm.* The Proportionality Norm says the following:

*The Proportionality Norm* (PN). A blamer's blame should be proportionate to a blamee's blameworthiness.

This seems especially apposite when it comes to partial excuses. Partial excuses make an agent less blameworthy than she would have been had she lacked the excuse. If an agent's blameworthiness is diminished, then both the manner and degree of our blame should be responsive to this fact. So, properly construed, PN seems to say the following:

*Retributive Proportionality Norm* (PN<sup>R</sup>). The manner and degree of blame should be proportionate to how blameworthy the blamee is.

Unfortunately, AOAA suggests that we likely fall short of PN<sup>R</sup>. In so far as we fail to properly attend to factors that partially excuse others from blame, it stands to reason that we'll tend to proceed as though they lacked any such mitigations. This in turn suggests that our blame will tend to be objectionable on *retributive* grounds—that we are liable to blame others more than is warranted. Consider what any one of us, qua onlooker in *Provocation 2* would likely be thinking to ourselves upon observing Jill's behaviour: 'What's *her*



problem?!' The issue is that we never really follow through on this thought. Just what *is* this person's problem? It's not unlikely that she does indeed have *some* problem if she's hurling shoes at people today! The issue here is procedural; we rarely even consider these questions. And inasmuch as we don't, we'll tend to leap to the conclusion that people like Jill are not even partially excused for what they do, and will fail to temper our blame accordingly.

It might have been expected that I'd be wrapping things up at this point. But there is a further complication that needs to be addressed. PN<sup>R</sup> may be the most natural way to interpret PN. But it isn't the only one. PN is in fact ambiguous—a point that often goes under-discussed, given the lack of forensic attention that's been devoted to it. There are other interpretations available, and it's worth asking whether we plausibly fall short of PN on these other readings as well.

When we speak of *degrees of blameworthiness*, we might (indeed, I suspect we often do) have in mind the disambiguation above, which modifies the attitudinal component: a partially excused agent is 'worthy of *less blame*' than others in the relevant comparison class (which may well include other possible versions of herself), who lack the partial excuse. But talk of degrees of blameworthiness can also be read as modifying the normative component; perhaps this agent is '*less worthy* of blame' than others.

This distinction is under-theorized in connection to the ethics of blame. But it is familiar in discussions of 'fitting-attitude analyses of value'.<sup>10</sup> Such approaches propose to analyse evaluative notions in terms of deontic ones. For something to be admirable, for instance, is for it to be fitting to admire it. Those attracted to such analyses are then tasked with accounting for *degrees* of value. And here too there is disagreement as to whether X's being more admirable than Y means that (i) X is more worthy of being admired than Y, or (ii) X is worthy of more admiration than Y (Rønnow-Rasmussen 2021, ch. 6).

There are reasons to be dissatisfied with (i), not least of which is that it fails to issue any ruling on the matter of *how much* we ought to admire something that's admirable or desire something that's desirable (Rabinowicz 2020, p. 14, cited in Rønnow-Rasmussen 2021, p. 107). And yet that approach still seems to be picking up on something important. Suppose both A and B achieve a failing grade

---

<sup>10</sup> Thanks to Wlodek Rabinowicz for raising this during a Q&A.

(below 50). Each may be equally worthy of a fail. But A (grade: 46) may nonetheless be *more worthy* of that fail than B (grade: 48) (Coates 2019, p. 238; compare Andersson and Werkmäster 2022, p. 198).

Can *both* approaches find a home in our account of degrees of blameworthiness? Coates (2019, p. 238) appears to think so. He interprets the property of ‘being worthy of more blame’ along the retributive lines suggested above, but takes the property of ‘being more worthy of blame’ to reflect there being stronger desert-based reasons to blame. On his analysis, ‘being more or less blameworthy’ is simply ‘a conjunction of these properties’ (2019, p. 242). Yet it’s not clear that such hybrid views can be made to work: for how exactly are these properties to be combined to yield a final judgement on degrees of blameworthiness? As Andersson and Werkmäster (2021) rightly ask, ‘Should the strengths be added together, multiplied, or is it a more complex function?’ It’s hard to see what could decide between these; ‘finding leverage for one function rather than another’ may well be ‘an impossible task’ (2021, p. 536). Perhaps that assessment is premature. But such issues should, I think, make us hesitant to stake our hopes (or indeed, our ethics of blame) upon any such combined analysis.

And so part of me is inclined to say that we had things right the first time: talk of degrees of blameworthiness is best interpreted as picking up on degrees of blame rather than degrees of worthiness—which in turn yields PN<sup>R</sup>. Nonetheless, there’s surely sense to be made of the idea that some people are *more worthy* of blame than others—and it seems to me that our ethics of blame ought to have something to say about that. Even if we want to resist building the notion of ‘more or less worthy of blame’ into our account of what it is to be ‘more or less blameworthy’, then, might it not have some other role to play? The answer that I now want to propose links degrees of worthiness to the *discretion* that we enjoy in how we *distribute* blame. The notion of worthiness that I invoke is admittedly peculiar, in that it is not purely desert-based (that side of things is already largely taken care of by PN<sup>R</sup>), but also incorporates reasons of other kinds. But it is, I think, a notion that will serve us well in building an additional dimension of proportionality into our ethics of blame.

To demonstrate what I have in mind, consider a modification of Coates’s scenario where the teacher receives permission to pass one more student—indeed, *any student*. Which student is *more worthy*

of this pass? If Coates has things right, then it's surely *B*; after all, there are stronger desert-based reasons to pass *B* than to pass *A*. And perhaps he does have things right if we insist upon the standard narrow construal of 'worthiness'. Yet to do so in this context would, I think, be to overlook the many other sorts of reasons that can rightly factor into such decisions. Suppose, for instance, that unlike *B*, *A* needs this pass in order to complete her degree. Forward-looking considerations such as these can, I submit, not merely be good reasons for passing *A* rather than *B*, but also make *A* a *worthier candidate* for a pass than *B*.

I admit that this stretches the meaning of 'worthiness'.<sup>11</sup> But it bears reminding that my project here concerns the *ethics of blame*, with the *justifiability* of *expressed* blame being my focal point. Forward-looking considerations may not ground blameworthiness per se. But they often prove important in justifying outward expressions of blame—or (as I'll now argue) in justifying discretionary choices in how we *distribute* blame among blameworthy agents.

Let me motivate these claims by way of example. In my department, a single person—call her 'Tammy'—is responsible for timetabling all classes. Each semester, there's inevitably a stuff-up despite my attempts at careful communication; for instance, a timetable that requires me to be in three places at once. It sometimes takes numerous email exchanges before I end up with a schedule that it's physically possible to adhere to. While I would never be outwardly hostile towards Tammy (or send her a passive aggressive email pointing out that unlike Hermione, I lack a time-turner and so cannot possibly keep to this agenda), I do grow angry and frustrated with this business. At this stage, my colleague—call him Craig—encourages me to *let it go*. Importantly, Craig doesn't deny that Tammy is somewhat blameworthy for causing this extra work. Nor does he suggest that I'm over-blaming her. (We both recognize that Tammy has a partial excuse; she is, after all, working with outdated timetabling software and has an army of angry academics breathing down her neck.) Craig's rationale is altogether different, and goes as follows: of all the things that one could get riled up about in this world, *timetabling errors* aren't particularly high up on the list.

---

<sup>11</sup> Though only slightly; it doesn't offend the ears to say that it is a *worthier use* of the pass to give it to *A* rather than to *B*.

As I interpret Craig, his point is that such errors aren't *an especially worthy target* of my blame—and I do think he is onto something here. But what is he onto exactly? One possibility is a concern about *diluting* blame; directing my blame towards partially excused mishaps such as Tammy's risks lessening its force whenever it is directed at real moral problems. Earning a reputation as a hothead could discourage others from taking me seriously on other occasions. Alternatively, Craig's concern may be one about *efficiency*: perhaps I only have so much moral energy in my stores, and this would be better spent on admonishing non-excused colleagues rather than partially excused ones. Or maybe my blame is unlikely to have much uptake with Tammy, who, given her circumstances, will (understandably) simply take me for yet another grumpy academic.<sup>12</sup>

Cases such as these suggest that an agent's partial excuse can render her *a less worthy candidate* for blame than others, for reasons that are partly forward-looking in character. Here, then, is the second proportionality norm that I want to propose:

*Distributive Proportionality Norm* (PN<sup>D</sup>). A blamer's distribution of blame should be sensitive to who is a more or less worthy candidate for it.

In my view, PN<sup>D</sup> reflects a bona fide norm on blaming—one we'd all do well to adhere to. Yet it seems to differ from its retributive counterpart in yielding only *imperfect* duties. Blaming others more than they deserve would seem to wrong them. But is it really wrong to distribute blame sub-optimally? Perhaps it would be inefficient to direct blame towards Tammy. But (provided that it is proportionate) this hardly seems wrong, and it certainly does not appear to wrong *her*. It is for this reason that we should, I think, view adherence to PN<sup>D</sup> as a good-making feature of blame, rather than a lack of adherence to it as a wrong-making one.

One potential concern for PN<sup>D</sup> is that it seems to conflict with another proposed norm on blame—in particular, with [Telech and Tierney's \(2019\)](#) 'comparative non-arbitrariness norm' (CNN), according to which it is *pro tanto* wrong to blame people differently for the same offence when there are no morally relevant differences between them. [Todd's \(2019, p. 369\)](#) example of blaming those

---

<sup>12</sup> Thanks to Edward Elliott for suggesting this third possibility.

involved in terrorist attacks in the West more than in the Middle East offers a nice illustration. So too does Telech and Tierney's own example of a parent who blames one of their children more than the other despite their misdemeanours being alike in all morally relevant respects (2019, pp. 27–8). We can add an example of our own into the mix here: suppose that I blame Tammy but withhold blame from other partially excused *male* colleagues like Tom. Telech and Tierney argue that such arbitrary differences in blame are morally objectionable. And that claim seems eminently plausible, on the face of it. How exactly do I propose to square this with my own claim that we're permitted to exercise discretion in how we distribute blame?

For my own part, I'm inclined to question whether CNN truly does reflect a norm on blaming. In the cases that motivate it, we appear to simply be adhering to a wider moral requirement *to be fair*. There's clearly unfairness involved in holding terrorists of a certain skin colour to different standards, or in being lenient towards one child and unforgiving towards another. But these are contexts in which someone is *already* required to be fair (with respect to blame as well as other things) and isn't. Compare the original timetabling case. Here, I'm effectively bestowing a kindness upon Tammy: I withhold blame because I recognize that it can be put to better use, and that Tammy, being partially excused, is not an especially worthy candidate for it. I wouldn't plausibly treat another partially excused colleague *unfairly* if I failed to extend a similar kindness to them.<sup>13</sup> Unlike our treatment of people of different races or different genders or our children, the distribution of such kindnesses isn't the sort of thing that's *supposed* to be fair.<sup>14</sup>

I've argued that given AOAA, we plausibly fall short of PN<sup>R</sup>. Is the same true of PN<sup>D</sup>? Or is this bifurcation of PN ultimately our salvation? Sadly not. Our susceptibility to AOAA suggests that we likely fall short of PN<sup>D</sup> as well. According to PN<sup>P</sup>, we ought to take excusing factors into consideration, and (at least sometimes) refrain from blaming those who aren't especially worthy of it. But in so far

<sup>13</sup> I follow Estlund (2008, pp. 67–8) in construing fairness as an 'occasional value'; something is unfair only when it is supposed to be fair but isn't. So something can be *not fair* without being *unfair*.

<sup>14</sup> CNN also seems to over-generate wrongful blame. Suppose I do withhold blame in Tammy's case. Given CNN, this morally constrains all of my future blaming interactions. Each time I encounter anyone else in a situation like Tammy's, I act *pro tanto* wrongly by failing to extend the same kindness to them.

as we overlook the circumstantial factors that partially excuse, we're unlikely to take them into consideration at all. Unlike Craig, I don't usually pause to consider the outdated software that Tammy is dealing with, or all of the other people bombarding her with complaints. In so far as I did, I'd be more likely to exercise discretion.

## VII

*Conclusion.* There's an old saying that 'you never know what someone is going through'. One often sees this accompanied by exhortations to be compassionate and understanding. It's easy to write off such talk as irritatingly sappy. But if the arguments of this paper work, then there may in fact be something to it. Even when we do have evidence that someone *is* going through something—or is at the mercy of something—our susceptibility to certain biases means that we're unlikely to recognize that evidence, or to properly attend to it or take it into account. Many of us seem all too ready shoot from the hip and assume the worst of other people. We plausibly owe it to others to give them a fairer hearing than this. Even when something resembling a hearing isn't feasible, we should at least be wary of over-confidence in our attributions of blameworthiness. Perhaps we should opt instead to give others the benefit of the doubt.<sup>15</sup>

*School of Philosophy, Religion and History of Science*  
*University of Leeds*  
*Botany House*  
*Woodhouse*  
*Leeds LS2 9JS*  
 UK  
[j.m.isserow@leeds.ac.uk](mailto:j.m.isserow@leeds.ac.uk)

---

<sup>15</sup> I am exceptionally grateful for all of the feedback that I've received on this piece of work. Thanks especially to Jessica Leech in her capacity as Editor, and to pre-readers and audiences at the Workshop on Epistemic and Practical Normativity at the University of St Andrews, the MetaFS Workshop at Frankfurt Business School, the Australasian Association of Philosophy Conference, the 'Standing to Blame (and Beyond)' Workshop at Lund University, and a workshop at Aarhus University. Thanks as well to Anna Stelle and Scott Sturgeon for all their work organizing and hosting the Aristotelian Society session.

## REFERENCES

- Andersson, Henrik, and Jakob Werkmäster 2021: 'How Valuable Is It?' *Journal of Value Inquiry*, 55(3), pp. 525–42.
- 2022: 'Normative Resilience'. *Utilitas*, 34(2), pp. 195–208.
- Austin, J. L. 1957: 'A Plea for Excuses'. *Proceedings of the Aristotelian Society*, 57(1), pp. 1–30.
- Baron, Marcia 2007: 'Excuses, Excuses'. *Criminal Law and Philosophy*, 1(1), pp. 21–39.
- Brink, David O. 2013: 'Situationism, Responsibility, and Fair Opportunity'. *Social Philosophy and Policy*, 30(1–2), pp. 121–49.
- and Dana Nelkin 2013: 'Fairness and the Architecture of Responsibility'. In David Shoemaker (ed.), *Oxford Studies in Agency and Responsibility, Volume 1*, pp. 284–314. Oxford: Oxford University Press.
- Coates, Justin D. 2019: 'Being More (or Less) Blameworthy'. *American Philosophical Quarterly*, 56(3), pp. 233–46.
- and Neal A. Tognazzini 2012: 'The Nature and Ethics of Blame'. *Philosophy Compass*, 7(3), pp. 197–207.
- Dover, Daniela 2019: 'The Walk and the Talk'. *Philosophical Review*, 128(4), pp. 387–422.
- Estlund, David 2008: *Democratic Authority: A Philosophical Framework*. Princeton, NJ and Oxford: Princeton University Press.
- Flick, Cassandra, and Kimberly Schweitzer 2021: 'Influence of the Fundamental Attribution Error on Perceptions of Blame and Negligence'. *Experimental Psychology*, 68(4), pp. 175–88.
- Gilbert, Daniel T., and Patrick S. Malone 1995: 'The Correspondence Bias'. *Psychological Bulletin*, 117(1), pp. 21–38.
- Holroyd, Jules 2021: 'Oppressive Praise'. *Feminist Philosophy Quarterly*, 7(4), pp. 1–26.
- Isserow, Jessica 2019: 'Moral Worth and Doing the Right Thing by Accident'. *Australasian Journal of Philosophy*, 97(2), pp. 251–64.
- 2022: 'Subjunctive Hypocrisy'. *Ergo*, 9(7), pp. 172–99.
- Jones, Edward E., and Victor A. Harris 1967: The Attribution of Attitudes. *Journal of Experimental Social Psychology*, 3(1), pp. 1–24.
- Kassin, Saul M., and Holly Sukel 1997: 'Coerced Confessions and the Jury: An Experimental Test of the "Harmless Error" Rule'. *Law and Human Behavior*, 21(1), pp. 27–46.
- Kelp, Christoph 2020: 'The Knowledge Norm of Blaming'. *Analysis*, 80(2), pp. 256–61.
- Littlejohn, Clayton 2020: 'Truth, Knowledge, and the Standard of Proof in Criminal Law'. *Synthese*, 197(12), pp. 5253–86.
- Malle, Bertram F. 2006: 'The Actor-Observer Asymmetry in Attribution: A (Surprising) Meta-Analysis'. *Psychological Bulletin*, 132(6), pp. 895–919.

- Piovarchy, Adam 2022: 'Situationism, Capacities and Culpability'. *Philosophical Studies*, 179(6), pp. 1997–2027.
- 2023: 'Situationism, Subjunctive Hypocrisy and Standing to Blame'. *Inquiry*, 66(4), pp. 514–38.
- Rabinowicz, Wlodek 2020: 'Are Probabilities Values?' In Ilkka Niiniluoto and Sami Pihlström (eds.), *Normativity: The 2019 Entretiens of Institut International de Philosophie*, pp.67–99. Helsinki: The Philosophical Society of Finland.
- Rønnow-Rasmussen, Toni 2021: *The Value Gap*. Oxford: Oxford University Press.
- Rosen, Gideon 2004: 'Skepticism About Moral Responsibility'. *Philosophical Perspectives*, 18, pp. 295–313.
- Royzman, Edward B., Kimberly W. Cassidy, and Jonathan Baron 2003: 'I Know, You Know: Epistemic Egocentrism in Children and Adults'. *Review of General Psychology*, 7(1), pp. 38–65.
- Sliwa, Paulina 2016: 'Moral Worth and Moral Knowledge'. *Philosophy and Phenomenological Research*, 93(2), pp. 393–418.
- 2019: 'The Power of Excuses'. *Philosophy and Public Affairs*, 47(1), pp. 37–71.
- 2020: 'Excuse Without Exculpation: The Case of Moral Ignorance'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 20*, pp. 72–95. Oxford: Oxford University Press.
- Smith, Angela M. 2007: 'On Being Responsible and Holding Responsible'. *Journal of Ethics*, 11(4), pp. 465–84.
- Strawson, P. F. 1962: 'Freedom and Resentment'. *Proceedings of the British Academy*, 48, pp. 1–25.
- Telech, Daniel, and Hannah Tierney 2019: 'The Comparative Nonarbitrariness Norm of Blame'. *Journal of Ethics and Social Philosophy*, 16(1), pp. 16–25.
- Todd, Patrick 2019: 'A Unified Account of the Moral Standing to Blame'. *Noûs*, 53(2), pp. 347–74.
- Wallace, R. Jay 1994: *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press.
- Watson, David 1982: 'The Actor and the Observer: How Are Their Perceptions of Causality Divergent?' *Psychological Bulletin*, 92(3), pp. 682–700.