Political Philosophy

Upadhyaya, Kartik. 2024. The Aptness of What We Do Together. *Political Philosophy*, 1(2): 545–574. DOI: https://doi.org/10.16995/pp.16914

Open Library of Humanities

The Aptness of What We Do Together

Kartik Upadhyaya, Philosophy, University of Warwick, UK, & Bioethics, Academy of Sciences, Czech Republic, kartik.upadhyaya.1@warwick.ac.uk

This essay outlines and responds to the problem of *appropriate piling-on*. Suppose that a person is responsible for having acted wrongly. It seems apt for you to blame that person. Now suppose that I also find out about the wrongdoing. If blame is an apt response for you, and there is no special difference between us, then it is surely an apt response for me as well. Imagine, however, that there are millions of people all in identical situations to you and me. It does *not* seem appropriate for all these people to blame the person together. But how can that be? Why can't we iterate the reasoning just outlined until we reach the umpteen-millionth person, and conclude that all of the millions can appropriately blame? I argue that everyone can blame inaptly even if each person blames aptly. The argument involves an account of when and why blame's outcomes make it inapt, drawing an unexplored contrast between *internal* and *external* standards of proportionality in blaming.

The Aptness of What We Do Together

Kartik Upadhyaya
Philosophy, University of Warwick, UK, & Bioethics, Academy of Sciences, Czech Republic

Writing in the *New Yorker*, D. T. Max narrates several episodes of what he calls 'The Public Shaming Pandemic'.[1] One features 51-year-old lawyer Lawrence Garbuz, who developed a cough and fever in February 2020. Initial tests suggested pneumonia, but his symptoms steadily worsened. Before being diagnosed with COVID-19, Garbuz attended a funeral and b'nai mitzvah, putting more than 100 families at risk of infection. At least 90 other cases were traced back to his. After New York's Mayor at the time, Bill de Blasio, had posted details about Garbuz on *Twitter* (now *X*), Garbuz's wife, Adina Lewis, responded with tweets in her husband's defence. This only expedited the snowball of online censure rolling Garbuz's way. One person tweeted, 'He continued to travel on metro north. He was coughing. His hands were filled with germs. Anyone he touched got sick … It was thoughtless and reckless.' Another said, 'I hope your business never rebounds for what your husband has brought upon us.' A third: 'He deserves to die. He's a scumbag. Endangered hundreds of thousands of people. He will never be able to live in New York again after this and he deserves it.' There were many hundreds more.

How blameworthy Garbuz is for what he did is difficult to discern from the evidence. But instances like his raise a normative problem. Suppose that a person has in fact acted wrongly and is doubtless blameworthy for doing so. It seems appropriate for you to blame the person. Now imagine that I also find out about the wrongdoing. We can presume that it is appropriate for me, also, to call the person out: if blame is an appropriate response for you, and there's no special difference between us, then it is, surely, an appropriate response for me. Thus, it seems appropriate for you and I both to blame. Now imagine that there are many millions of people, all of whom are in similar situations to you and me – all of them have discovered the wrongdoing. It does *not* seem appropriate for all these billions to blame the person together. But how can that be? As I suggested, if blaming is appropriate for me to do, and appropriate for you to do, it is appropriate for you and I both to do. Add a similarly situated third, fourth, and fifth person: why wouldn't blaming be appropriate for all five of us to do? Why can we not iterate this reasoning for each additional blamer until the umpteen-millionth, and conclude, on that basis, that all of the millions can appropriately blame?

Call this the problem of *appropriate piling-on*, which Section I outlines in more detail. It is a pressing problem to try to make sense of. Social media technologies have

[1] Max 2020.

increased the possibilities of ever larger pile-ons. No doubt we've long practiced public condemnation offline, too – think e.g. of town folk hurling groceries at transgressors locked in the local stocks. But the emergence of online communicative fora, and the pervasive role these now play in the social and political lives of billions the world over, has taken things to a new scale. More people can acquire information about allegations of immorality, and people can acquire that information more quickly and easily. Users, likewise, can give a louder public voice to their judgments at lower personal cost. It is possible now to formulate and express our thoughts unburdened by social anxieties and other difficulties that come with participation in more traditional discursive platforms, such as the town hall, or *Question Time*. What people say on social media also has longer and wider reach. We can directly address people who we may never otherwise encounter in our lives, with nation-sized audiences who may otherwise pay us little attention. Many of those audiences are able to share that content with still wider audiences of their own.

It is tempting to think that this is all for the better. Public discourse and mutual accountability are good things. Social media platforms facilitate more of those good things, and more publicly at that. However, sometimes we can have too much of a good thing. I will argue that mass piling-on can be collectively inappropriate, in the following sense: even if each participant does something apt, everyone can together do something inapt. Acting together is inapt when and because the harm caused by the overall practice is unjustifiable relative to the proper aims of blame. I will call this thesis the *Proportionality View*. Section II introduces this view, showing how it explains contrasting judgments about smaller and bigger numbers of blamers.

As mentioned, this response to the problem divorces two issues: first, the aptness of what each person does individually; second, the aptness of what we do together. Some might reject that separation. According to the *Inheritance View*, if and when everyone's piling-on is inapt, that collective fault is inherent in each person's involvement. According to the *Composition View*, if and when each person behaves aptly, it must also be the case that the set composed of all these people behaves aptly. According to the *Equivocation View*, differences in individual and collective aptness trade, illicitly, on an ambiguity between different types of blaming: *communicative* and *non-communicative*. Section III argues against all three alternatives. Because of the distinct ways in which blaming, both communicative and non-communicative, does harm as a *collective* practice, the aptness of each person's conduct is rightly divorced from the aptness of everyone's conduct.

It is unusual to claim, as the Proportionality View does, that blame can be rendered either apt or inapt by instrumental facts, including its harms. I account for this quirk by arguing that the aptness of blaming is determined by how well it instantiates a

specific subset of values: values that play a special role in vindicating our accountability practices. Piling-on is inapt when its outcomes cannot be vindicated by those specific values. Section IV presents this argument for the Proportionality View – I call it the *Internal Values Account* – and outlines some of its virtues. Section V explains how the account is able to differentiate the claim that blaming is apt from the claim that it is a good or bad idea in view of its outcomes.[2]

Proportionality standards are part of legal tradition and have purchase in framing content moderation policy.[3] They are also established in philosophical debates about content moderation[4] and public shaming[5]. What's at stake, normatively speaking, between those discussions and the contention of this essay – that piling-on is objectionable because it is *inapt*?[6] Section VI explores a distinction between what I call *internal* and *external proportionality*, and compares that distinction with Paul Billingham and Tom Parr's application of *narrow* and *wide proportionality* to the context of public shaming.[7]

Before we get going, a confession. I will make lots of assumptions below about *blaming*: about the diverse activities that belong to blaming practices, the value that these activities bring to moral life, the burdens with which they saddle wrongdoers, and why the number of people involved in blaming should make us worry about the appropriateness of imposing those burdens. Whilst my responses to those questions are designed to be plausible and attract some consensus, it's difficult to please everyone about these matters, and I shirk the task of offering a complete account of all the assumptions made.[8] On its most ecumenical reading, then, the argument to come might be understood as conditional on the picture of accountability practices we will be drawing up. If that picture turns out reasonable on the whole, a problem looms with the structure that I outline next, as do the ingredients for a solution. Might we get similar results on radically different pictures of our blaming practices? I do suspect so, but will tend to defer that question to future work.

---

[2] I thank an anonymous referee for this phrasing.

[3] It is common among the recommendations issued by *Meta*'s new *Oversight Board*, and is invoked more than 73 times in the UK Government's 2020 *Online Harms White Paper*.

[4] Howard 2019; forthcoming.

[5] Thomason 2021; Fritz 2021; Frye 2021; Billingham & Parr 2020.

[6] Fritz (2021, p. 690) remarks in passing that 'the effects of our disapproval are often more intense than is appropriate'. This essay might be read as one way to elucidate and account for Fritz's verdict.

[7] Billingham & Parr 2020.

[8] Influencers of the picture I will offer include: Bennett 2002; Pickard 2011; McKenna 2012a; McGeer 2019; Fricker 2016; Sliwa 2020.

I. THE PROBLEM

The problem of appropriate piling-on is a trilemma of claims, all of which seem attractive at first, but cannot be jointly true. Here are the three claims:

i. It is appropriate for each person (identically situated) among many to blame a person who is responsible for wrongdoing.

ii. If it is appropriate for each person (identically situated) among many to blame a wrongdoer, then it is appropriate for everyone to blame that wrongdoer, regardless of how many blamers there are.

iii. It is not appropriate for everyone (identically situated) to blame a wrongdoer if the number of blamers is sufficiently high.

Let me explain a bit further the initial attraction in each idea.

I take it that *i* is the least resistible claim. What counts as a blaming response is, as I just highlighted, a vexed issue, and we will survey a few possibilities shortly. But suffice to say for now: very few deny that at least some form of blaming response to wrongdoing is appropriate.[9] Some theorists go so far as to argue that blame's being appropriate is part of what constitutes an act of wrongdoing.[10] Whether or not we ought to go that far, the fact that people take that view seriously indicates a deep relationship between wrongdoing and the aptness of blaming.[11]

But why believe that claim *ii* is true? How might the idea that blame is apt for each person support the idea that it is apt for the collective, regardless of how many people in that collective? It is not in general odd to think that other people doing something can affect the morality of my doing the same. The fact that others will rescue a drowning swimmer might make my rescuing the swimmer less urgent. The fact that others contribute to climate change might make my own contribution more urgent. The fact

[9] Those who doubt blaming practices are a good thing tend to think that hostile attitudes are essential to blaming: see e.g. Kelly 2018; Pereboom 2014; Pickard 2011.

[10] Compare D'Arms & Jacobson 2003, Wallace 1994, and (arguably) Strawson 1962.

[11] Some philosophers think that appropriate blaming has a 'business condition': the blamer must bear an answerability relationship with the blame which makes the blamee's conduct 'the blamer's business' (Duff 2011; Radzik 2011; Seim 2019; Upadhyaya 2024). Those drawn to this idea might doubt *i* on the grounds that it is not 'the wider public's business' to call a person to account for many types of wrongdoing they might commit (wrongs of betrayal between intimates might be central among these). For this essay's purposes, let us assume any applicable business condition is met. Assume there is an obvious public interest in the type of wrongdoing that is subject of piling-on (e.g. wrongly spreading a virus on the streets).

that others have tidied the kitchen might make it more incumbent on me to tidy the kitchen. And so on.

But the aptness of reacting to wrongdoing seems to work differently. The aptness of one person's response does not normally change just because others respond the same way. Suppose, for example, that Zara is responsible for wrongdoing, and that Ali blames Zara. By itself, this makes no difference to whether it is apt that Bea blames Zara. Assuming that Ali and Bea are identically situated, the fact that another blamer exists does not seem to alter what is apt for Bea to do. Call this feature of aptness *addition-resistance*.

Addition-resistance suggests that the same reasoning applies to each additional blamer. Adding a third identical blamer makes no difference to our evaluation of the first two, or to all three blaming. Adding a fourth makes no difference to our evaluation of the first three, or to all four blaming. The natural conclusion: adding an *n*th identical blamer makes no difference to our evaluation of *n*-1 blamers, or to all of *n*'s blaming. After all, these people are all responding to the same facts in just the same way. It is unclear why what is true of additions to small values of *n* would not also be true of additions to big values of *n*.

However – and this is the problem – claim *iii* calls that reasoning into question. When the numbers are immense, e.g. in the millions or billions, I doubt that many would consider the practice appropriate. Intuitively we do distinguish additions to small groups of blamers from additions to massive groups (and perhaps equivalently, between small additions and massive additions).

So what gives? Must we abandon that intuition about very big numbers (*iii*)? Must we abandon addition-resistance (*ii*)? Or, must we abandon the idea, at least when it comes to mass pile-ons, that blaming is an apt response for each person to begin with (*i*)?

Before outlining this essay's response, I should call to attention the distinctiveness of this problem. Much philosophical discussion is had about situations in which individual contributions to a series of acts seem to differ in their moral valence from the series of acts as a whole. But our task is not as simple as exporting conclusions from those discussions to the case of piling-on. The latter has special features.

One puzzle in the ethics of killing may appear analogous. Millions of armed assassins queue up to wrongly kill an innocent person.[12] The target seems justified in killing, in defence of their life, each wrongful aggressor appearing at the front of the

[12] McMahan 2017.

queue – each individual defensive act seems objectionless. Taken together, however, these defensive acts result in millions of deaths for the sake of a single life – the series as a whole looks objectionable. But unlike piling-on, the queue of aggressors is a *one-to-many* situation. In these situations, the series of acts (defensive killings) that seem individually justified are performed by a single agent (the innocent target), affecting multiple patients (the many aggressors). Appropriate piling-on is, conversely, a *many-to-one* situation. Multiple agents (the many blamers) do the act that seems appropriate (blaming), affecting a single patient (the person blamed, hereafter the blamee).

The reason appropriate piling-on is distinctive is not merely that it has a many-to-one structure. A famous set of examples mirroring that structure, Derek Parfit's *torturers*, involves many people raising the current in an electrical torture device. Suppose the device is hooked up to a single innocent person. Each person presses a button in order to raise the current by a tiny, imperceptible amount. All of those people pressing together results in the victim's agonising torture.[13] But this many-to-one case is also disanalogous. The difficulty with the harmless torturers lies in explaining the idea that each torturer acts wrongly, despite having an imperceptible effect on the victim. In the case of appropriate piling-on, it may be true that each blamer's contribution is imperceptible, given that so many others are contributing in the same way. But I take it one is less confident that each blamer acts wrongly. As I said, we might rather think that in holding someone accountable with accuracy, each blamer does something positively *apt*. Our task is to explain the different reaction to many people doing the same seemingly apt thing.

Once more, the obvious route out of the problem is to abandon at least one of its constituent claims (*i-iii*). In what follows, I reject claim *ii*. I will offer an explanation of why it is that the aptness of each person's blaming needn't go hand-in-hand with the aptness of everyone's blaming.

II. THE PROPORTIONALITY VIEW

My argument against *ii* is based on the *Proportionality View* introduced at the outset: the view that blaming a wrongdoer is inappropriate when and because the harms it involves are unjustified relative to the value of blaming that person. Let me begin by clarifying why this view is grounds for rejecting *ii*.

[13] Parfit 1984, pp. 80–82. For discussion of this case in relation to online criticism, see Billingham & Parr 2020.

Suppose, as will be argued later, that for a given act of culpable wrongdoing, there is a maximum amount of harm one can appropriately inflict on the wrongdoer by blaming them. High numbers of people blaming can together harm that wrongdoer in amounts that are unjustified relative to the value of blaming that person at all. If those amounts are reached, the practice as a whole is inapt, according to the Proportionality View. However, in those same situations, it is also possible that *each* person's individual blaming act does *not* involve a level of harm that would be unjustified relative to the value of blaming the wrongdoer. So, each person's blaming may not be inapt in that sense. It is therefore possible both that each 'piler-on' blames aptly and that all 'pilers-on' together blame inaptly. And that is the possibility that claim *ii* rules out.

Notice that judgments about appropriateness in these contexts seem to correspond to the amount of harm one would reasonably expect to befall the blamee. Let us start, again, with a single blamer: suppose Ali accurately blames Zara in response to Zara's culpable wrongdoing. Bracketing the blame sceptics set aside earlier, most are likely to think that Ali behaves aptly. However, this is not because we imagine Ali's blame to be harmless. To the contrary, many advocates of blame point out how it carries a 'sting'.[14] The point is that the sting of Ali's blame can, in some way or another (more on possible ways later), be justified in the light of whatever accounts for the value of blaming Zara.

Now add the second blamer, Bea, to the set. The addition of Bea is not of major concern even if we would expect Bea's blame to come at some additional harm to Zara. Again, we might be able to explain this by appeal to the value of blame – the value of blame is such that it can often justify harm done by adding a second blamer.

Now add a thousand blamers. At this point, one pauses before judging the thousandth blaming act appropriate. The reason I at least pause, though, is not necessarily because thousands of people blaming Zara is entirely devoid of value. Rather, the value achieved by these thousands of acts fails to justify the harm I would expect to be at stake.

To make this less abstract, think about some concrete harms and values. Often accompanying the initial shock or 'sting' of hostility and judgment are further distresses such as guilt, embarrassment, anger at, and disappointment in, oneself. Other harms, salient on online platforms, include social disesteem or exclusion by others. Disesteem and exclusion can result in downstream burdens, but their concomitant harms can also be more direct: consider anxiety about the possibilities of social exclusion or loss of esteem, and about whether one is the kind of person who merits these negative social reactions.

[14] See Pickard 2013, p. 619, and Fricker 2016.

Turning to values, one set widely associated with blaming are motivations in the blamee to remedy their fault based on a moral understanding of that fault.[15] Another set of values, salient in public fora, include the expression and display of respect and solidarity towards victims of wrongdoing.[16] A third involves victims themselves exercising a voice, having a platform publicly to share their account of moral injury, and to demand the responses that wrongdoers owe them.[17]

Those are weighty goods! They certainly seem worth achieving at significant cost. Suppose, then, that respect and solidarity towards victims, or motivation to make amends and reform, inevitably come at non-trivial levels of guilt, anxiety, and loss of face. That may be a price worth paying.[18] This, I think, is a key reason why the harms expected of paradigm one-to-one blaming interactions do not threaten the aptness of those interactions. There may be other reasons too: some believe that the harms of blaming are themselves valuable, e.g. because Zara deserves blame as someone responsible for wrongdoing.[19] Others believe that harms are a necessary part of something valuable, e.g. that guilt is necessary for moral understanding, or social exclusion is necessary to show respect for victims.[20] But these are routes to the same conclusion: the values of blaming are standardly able to justify the harm it does.

Furthermore, those same values are also able to justify multiple blaming responses to the same instance of wrongdoing. To be sure – when more people blame, the harm requiring justification tends to increase. Not only can additional blaming interactions each contribute their own 'sting', as discrete psychologically distressing experiences; they can also compound the harms of a single interaction. When many rather than few blame, threats to social standing, for example, will seem and become more real. At the same time, however, the addition of blamers may enhance some of the values outlined above. The fact that multiple people confront a wrongdoer about their wrongdoing can underscore the seriousness of that message to the wrongdoer. Multiple exhortations and reminders of a person's need to reflect on their behaviour and respond accordingly

---

[15] See footnote 8.

[16] Billingham & Parr 2020.

[17] See Upadhyaya 2024; also compare Radzik 2011, esp. p. 571.

[18] Fritz (2021) is sceptical about the worth of a kindred trade-off in the context of online shaming. Fritz weighs '*openness*' – i.e. limitless speech – against '*decency*' – i.e. inclusion within a community. While I tend to agree with Fritz's argument, it differs in scope from mine, as openness is more extensive than the goods of accountability listed here. The goods of accountability are consistent with more restricted speech.

[19] Nelkin 2016.

[20] E.g. Fricker 2016; McKenna 2012a; Bennett 2002.

can provide more powerful motivation than a single interaction. We might also think that solidarity to victims is better communicated when several members of a community speak out.

Albeit enhanced in these ways through multiple interactions, the values just described are hardly *sacrosanct*. There are, surely, restrictions on our ability to appeal to them to justify a practice we engage in.

Consider four restrictions familiar to wider ethics of harm, but also to the ethics of public shaming.[21] First, the blamee might show that they have gained pertinent moral insights and motivation, for example through recompense and willing dialogue with those calling them to account. As the values sought have already been secured, they do not justify additional harmful blaming. Additional harms are in this case *gratuitous*. A good example is the old saying – 'Once a meme, always a meme.' A 'memed' person continues to face distressing critical content no matter how much apology, reconciliation, reform, and so on, has taken place.[22]

In other situations, the harm done might be *counterproductive*. The blamee may become too intimidated or traumatised to absorb the message, or to engage in the moral self-reflection required of them. Self-harm in response to online shaming is a more extreme though equally recognisable instance.

Third, the burdens may be of such a magnitude that reform and respect for victims cannot support their imposition. Here, the practice is straightforwardly *disproportionate*. Long-term harm to mental well-being, including self-harm, are often disproportionate in this way, as are some levels of socioeconomic damage. (Confronting someone about their minor breach of etiquette isn't worth blacklisting them from future job prospects.)

A blaming practice is not off the hook even if it manages to be productive in achieving worthwhile goals, and even if it does not involve pointless further intervention once achieved; for those same goals might have been achieved in less destructive ways. Suppose that in order to understand the seriousness of their offence, it is sufficient for a culprit to face the consequence that others will refuse to support them until they apologise and explain themselves. Additionally exposing them to unsettling threats of violence is unjustified as *unnecessary* harm.[23]

---

[21] These standards are part of Billingham & Parr's (2020) framework of constraints on shaming; I later tease out some differences in the view I am defending here (Section VI). (For a more complete taxonomy of proportionality standards, see Tomlin forthcoming.)

[22] Compare Ellen's (2023) headline.

[23] Why do these phenomena punctuate online discourse? One possible set of factors suggested earlier, again explored by Billingham & Parr 2020, is the size and reach of social media com-

The Proportionality View holds that if any such objection applies – if its harms are not justified relative to its value – blaming is, for that reason, inapt. I have just been illustrating why this view is grounds for rejecting *claim ii* in the light of how the addition of blamers affects both harm and value. In the light of those effects, proportionality restrictions may apply to harm caused by multiple rather than single contributions. Adios addition-resistance.[24]

The foregoing is only a sketch of the thesis, to be developed further below. Natural question-marks will arise about how to explicate the proposed relationship between aptness, on the one hand, and the harms or goods of blaming on the other. For example, some would argue that it can be apt to blame harmfully even if this does not secure reform, as when we blame psychopathic or other morally hardened people incapable of reform. Others might argue that it can be inapt to blame harmfully when this does secure reform: it seems inapt, for example, to pile more harm on a wrongdoer to incentivise lots of *other* people's reform.[25] How, more generally, are we to separate the question whether blaming is apt, or inapt, from whether its outcomes reveal it to be a good or bad idea?

We will revisit these concerns in earnest later on.[26] Another question mark to be raised about the Proportionality View is that it is not the only possible account of our problem. Some may be instead inclined to reject *claim i*, to reject *claim iii*, or indeed to reject *claim ii* on grounds other than those adduced here. Let us next consider these alternative strategies.

---

munities. See also Frye 2021: despite framing the problem of scale as distinct from proportionality, the social costs of shaming Frye outlines seem subsumable under the proportionality standards outlined here. For example, the bigger the community of critics, the stronger the reasons that targets of criticism have to worry about their reputational futures. (More on that point below in Section III.C.)

Secondly, online commentators are not always confronted with the plight or identity of those they harm; it may be psychologically easier to harm those from whom we are distanced in these ways. In addition, social media platforms lack a history of established accountability mechanisms. (See Thomason 2017 and Norlock 2017.) Users can bypass penalties or censure through acting under false identity, deleting, and blocking accounts. This may diminish motivation to consider the harmful implications of posting, and to refrain from posting even if the user does consider those implications.

[24] More pedantically: addition-resistance needs revising so as to be made conditional on a certain kind of proportionality standard (still to be explicated in full).

[25] I owe these instructive examples to the Editor.

[26] In Section V.

III. AGAINST THREE ALTERNATIVES

A. The Inheritance View

One counterargument to the view just presented runs as follows. Piling-on is a bad practice. Since the practice is bad, and since it is constituted by the actions of those participating in it, each participant (to the extent that they're responsible for their blaming decisions) must be culpable in some way, or at least cannot appropriately take part. Call this the *Inheritance View*. It holds that when *claim iii* is true – when piling-on is inappropriate overall – *claim i* must be false. Each piler-on inherits whatever makes the overall practice inappropriate by responsibly involving themselves therein. The solution is therefore to reject *claim i*.

Views with a similar structure enjoy some support among philosophers of collective action. Some argue that participants in collective harm can be open to objection while contributing little or nothing to that harm, in virtue of their attitudes towards the possibility of inflicting it; attitudes displayed through their participatory acts.[27] This is no doubt plausible when it comes to real-world practices. There are so many deficiencies displayed in how we blame each other; no less do these glare in online pile-ons. Consider first deficiencies in the accuracy of blame's content, as seen in Garbuz's opening example. It is one thing to exclaim that someone acted objectionably, and to insist they must address that fact. It is quite another to say that they should therefore be excommunicated. Or have their livelihood destroyed. Or die! Such remarks are false if not absurd. False blame is inapt.

I also mentioned that while evidence on Garbuz's blameworthiness seemed inconclusive, people jumped on the pile regardless. That is another familiar phenomenon. Though social media is also a means of sharing a good deal of sound information, many are uncritical in lapping-up the torrents of bad evidence in circulation. Technological advances have also diversified bad evidence. Agents of *disinformation* obscure both the sources of information shared and the identity and motives of those sharing it. *Deepfakers* manufacture material that appears to have 'smoking gun' status. In short, much blaming that goes on is unwarranted by the quality of the evidence available. Unwarranted blame is inapt.

Other dimensions of inapt online criticism are motivational. Why, some might wonder, would anyone knowingly jump onto mass piles? People who make false or exaggerated claims, and who are quick to accept evidence of wrongdoing, are often

[27] Compare e.g.: Wieland 2022; Driver 2015; and Kutz 2000. (Against this idea, see esp. Nefsky 2023.)

angry people who need scapegoats. Some are grandstanders, whose main intentions are to project to the world that they lie on the right side of morality.[28] And some are *trolls* who jump on the bandwagon just for kicks. Blaming for such reasons – because one is generally angry, or grandstanding, or thrill-seeking – falls short of blaming out of genuine commitment to the moral values at stake. It seems inapt to hold others accountable to values without ourselves being committed to those values.[29]

These are real and legitimate objections to many people's involvement in piling-on. But as contingent facts about the real world, they do not show that the inaptness of piling-on must be inherent in the conduct of each participant. For our purposes, we can simply stipulate that evidence of someone's wrongdoing is clear, compelling, and understood by participants, all of whom are sincere in wanting to be part of an accurate process of calling the wrongdoer to account for themselves. In this instance, I see no reason to suppose that any single blamer behaves inaptly. Still this is a possible instance of inapt piling-on.

B. The Composition View

Some may have doubts about the point just made. Perhaps people only consider pile-ons inapt because when imagining them, they imagine real people blaming in the unattractive ways that real people do. However, once we assume *claim i* – once we bear in mind that each piler-on is well-motivated and epistemically responsive – claim *iii* no longer seems true. Call this response the *Composition View*. It says that the aptness of a practice is the sum of its parts: a practice is apt overall when composed of, and only of, apt behaviour.

The Composition View overlooks the difference between intending harm and contributing to it.[30] Calling people out hurts them even if hostility is no part of our motivation; this should be recognisable enough from experience of one-to-one blaming interactions. It may be possible sometimes, or possible in principle, to face one's moral failings, and others highlighting of it, entirely free of distress;[31] but whether or not they are possible to avoid, guilt, shame, anxiety at the prospect of social exclusion and the other harms mentioned earlier, are extremely difficult to avoid – more so under the

---

[28] Tosi and Warmke (2020) discuss both piling-on and exaggeration as types of grandstanding. As explained below, piling-on as understood in this essay needn't involve grandstanding.

[29] Compare Shoemaker & Vargas 2019 with Rossi 2018 and Todd 2017.

[30] Compare Fritz 2021, p. 498.

[31] For arguments against this possibility, see McKenna 2012b.

scrutiny of big numbers. The motives and warrant of any blamer underdetermine the harm blame does, and that harm, I maintain, can make the practice inapt.

Try to imagine that this is not the case. Imagine *proportionate piling-on*, in which everyone holds a wrongdoer accountable with accuracy and good faith, together inflicting no more harm than is justifiable, no matter how many blamers are involved. Here I agree that the practice is apt. But this further supports the idea that concerns about aptness depend on the expectation of unjustified harm. The Proportionality View explains why it's in those cases that the problem of appropriate piling-on arises at all.

The Proportionality View also explains why there is no symmetrical problem of *appropriate piling-on of praise*.[32] Trillions of people accurately praising someone who is responsible for heroism does not strike one as inapt. Again, we can explain this by pointing out that additional praise does not heighten expectation of unjustified harm. And that seems to me the right explanation.

At this juncture, proponents of the other view rejected above – the Inheritance View – may be minded to press the following objection.[33] I have just said that the problem arises only if harm enters the picture, intentionally or otherwise. However, harms that are unintentional might be foreseeable. It might be reasonable for each participant in piling-on to infer a risk of disproportionate outcomes. Is it not then inappropriate for each of them to take that risk?

On this argument for the Inheritance View, what's inherited in the actions of each person is the risk of harm done by a sufficient number, rather than the harm in fact done by everyone. Each person's decision to risk causing harm that exceeds some proportionality threshold is inaptly negligent or reckless. Shelly Kagan makes a comparable argument about numerous collective action dilemmas.[34] He emphasises the chance that each person has of being what he calls a 'triggering agent' – the agent that ends up bringing a situation to an intolerable level. Imagine a chicken shop that orders chickens from a farm in batches of 25. They do this in direct response to customer demand: every 25th order of current stock triggers the shop to request another 25 chickens to be slaughtered at the farm. Then, a customer whose individual order brings the total number to a sufficient level (say, every 25th customer) triggers a very bad outcome (25 chickens dying for the sake of that individual's consumption). Any given customer has a chance of being that triggering agent. Of course that chance is small for

[32] Thanks to Victor Tadros for discussion.

[33] Thanks to Massimo Renzo for discussion.

[34] Kagan 2011, pp. 119–129.

each of them.[35] But, Kagan argues, since the badness of 25 chickens dying is major, and the good obtained minor, the disvalue of that possibility is high enough to overwhelm its remoteness, rendering negative the expected value of each purchase.[36] One may argue by analogy in the case of piling-on: unjust harm is also a serious disvalue, and one that I have claimed may be triggered by a sufficient number of blamers. Although the chance of each piler-on's triggering that possibility is small, its disvalue is high enough to render negative the expected value of each contribution.

But I doubt the analogy quite works. True, unjust harm is very bad. However, a wrongdoer recognising the seriousness of their misdeed, and being accountable to others for it, is very good. (Much better than chicken!) Each piler-on has a chance of triggering that good. Each thus has a chance of doing something very good as well as something very bad. It is unclear, therefore, that each piler-on should expect the value of their contribution to be negative.

C. The Equivocation View

Even if the Proportionality View is right to take issue with claim *ii* rather than *i* or *iii*, perhaps there are better arguments against *ii* than the view offers. On one version of this third possibility, the problem of appropriate piling-on equivocates between different types of blaming activity.[37]

Suppose first that each piler-on blames in a strictly *non-communicative* way – each adopts morally accurate attitudes about the wrongdoing without communicating those attitudes to the wrongdoer. In this case, claim *i* and *ii* are plausible, yet claim *iii* is no longer. There is no concern about an unlimited number of people together adopting blaming attitudes. So the trilemma dissolves.

Next suppose that pilers-on blame the wrongdoer in a *communicative* way – each communicates their blaming attitudes to the wrongdoer. While *iii* might then be plausible, since there *is* concern about unrestricted communication, *iii* needn't be in tension with *i* or *ii*. For *i* remains plausible in reference to *non*-communicative blaming: it remains plausible that each person behaves aptly inasmuch as each adopts the right

---

[35] 1 in 25.

[36] The expected disvalue of any given purchase is the disvalue of 25 chickens dying multiplied by that event's probability (1/25): i.e. one chicken's death. The expected value of any given purchase is one person's pleasure in consumption, which is much less important. That is why any given purchase's net expected value is negative.

[37] This excellent idea I owe to an anonymous referee.

attitudes, yet there is nothing inconsistent about believing this whilst also believing that all of them behave inaptly in doing something *else*; namely communicating those attitudes. The trilemma again dissolves.

Diagnosis? According to the *Equivocation View*, claim *ii* generates the problem through sleight of hand. Its consequent draws a false inference about communicative blaming (all blame inaptly), from an antecedent about non-communicative blaming (each blames aptly).

This diagnosis appeals in that it allows us both to defuse the trilemma – to view *i*-*iii* as jointly consistent – and to maintain the instinctive objection to the appropriateness of piling-on. Yet we cannot so easily have our cake and eat it. The problem persists if we clear up these ambiguities and ensure that *i* through *iii* refer to blaming activity of the same type.

Start with the case of communicative piling-on. Recall the stipulation above that each person involved is well-motivated, seeking to engage the wrongdoer in response to an accurate understanding of the gravity of their wrongdoing. True enough, it is consistent with the idea that all these people communicate their attitudes inaptly, that each of them blames aptly inasmuch as each adopts the right attitudes. However, I maintain that beyond adopting apt attitudes about the wrongdoer's conduct, each person can also aptly communicate these in a call for the wrongdoer to account for that conduct. I will give more explanation below, but it is at least arguable that true attitudes about a person's responsibility for wrongdoing are aptly communicable to that person.

To clarify, the presumption here is not that it is always apt to communicate true attitudes about wrongdoing. Suppose Ali learns of Zara's wrongdoing, signs into a social media forum and can clearly see that Zara is already trending as the topic of a rabid dogpile. It would be inapt for Ali to add insult to injury by communicating true blaming attitudes to Zara.[38] In this context, adding one's penny's worth is closer to Kagan's chicken consumer. Ali's expectation about the value of their contribution ought to be 'nil at best'; the goal of achieving accurate accountability cannot then vindicate what Ali does.

However, in other contexts, blamers-in-prospect are not so clear about whether an instance of wrongdoing will be suitably addressed, and at what cost. These issues are often unclear in nascent stages of an online pile-on. Contributions in nascent stages may be made synchronically, or moments apart, affording little time for any agent to form well-founded expectations about how others are acting. In these more uncertain situations, each well-motivated agent might be apt to seek accurate communication with the wrongdoer, even if all such agents fail to act appropriately together.

[38] This is one way to understand Bhargava's (2020, pp. 387–390) 'Dependence Principle'.

Turn now to the trickier case of *non-communicative* piling-on. The Equivocation View is right inasmuch as this is not the central example of our problem. The reason online discursive media make the problem so palpable has of course to do with communication's greater capacities for harm. Nevertheless, the Equivocation View understates the perils of unrestricted non-communicative blaming.

Here is a live possibility. We have interests in other people not forming negative attitudes about us even when these people do not communicate those attitudes to us. Blaming attitudes are among the negative attitudes which set back our interests in that way.

I earlier alluded to one instance of this setback – when a person suffers distress in thinking or realising that people now see them in a worse light. Often non-communicative blaming won't engage *that* interest, of course, as difficult to know that people are seeing you in a worse light unless those people *tell* you so. Notice, though, that this difficulty is not a blanket impossibility. Non-blaming witnesses might relay their gossip. It might even be possible to know what is going on without anyone communicating with you. Maybe some among us (think of the paranoid or hyper-observant) have special antennae for blaming practices developing around them. Maybe some do not need telling.

But, in my view, a further underlying reason why a person suffers distress in coming to know about others' attitudes is that they come to know that something independently bad for them is happening: namely *that people are seeing them in a worse light.* This fact seems independently bad for us in various *instrumental* ways. Blaming attitudes can influence behaviour other than by moving a person to reveal them to their target. For instance, these attitudes can move people to shun or distance themselves from a wrongdoer, and might therefore contribute to some of the more familiar harms of blame (ostracism, socioeconomic disadvantage, etc.) discussed in the previous section.

Arguably, blaming attitudes about us might also be *non-instrumentally* bad for us. Jay Wallace writes – 'we do not wish to inhabit a social world in which such attitudes are harboured toward us, regardless of whether they are expressed or experienced by us. It is bad for us, in its own right, to be the target of such attitudes'.[39] Is Wallace's view compelling in our problem case: when the attitudes harboured towards us respond to the truth about what we've done?[40] That question pulls me in different directions,

[39] Wallace 2019, p. 92. Compare also Fricker 2016 and Slavny 2017.

[40] An interesting approach to this question is to think about people's interests in unconditioned social interactions (see Pressly 2024).

but whatever the answer, blaming attitudes seem capable of some communication-independent harms.

What is more, the harms of being seen in a worse light might be inappropriate to inflict in the collective sense described by the Proportionality View. (Two or three New Yorkers' thoughts about Gabruz's virus protocol do little to make Gabruz worse off, but hundreds or more having those thoughts might do a great deal.) The bigger the number of people who adopt blaming attitudes, the more serious their harms. Like any harm, these non-communicative harms may or may not be vindicated by the value of blaming the person to begin with.

IV. APTNESS AS INTERNAL VALUE

So far, I have outlined how the Proportionality View explains the problem of appropriate piling-on and argued that it does so more appealingly than some other explanations. But what explains the view itself? Why surmise, in the first place, that blame can be rendered inapt by the harm it does? To answer this, I now offer a general account of what determines blame's aptness, and explain some of its virtues.[41]

A. Internal Vindication

The main thought in the general account is this: How appropriate an activity (*a*) is, is set by the values of a practice (*p*) of which it forms a part. *a* is appropriate to the extent that it realises the values of *p*, and inappropriate to the extent that it fails to realise the values of *p*. The values to underscore here are the values <u>of *p*</u>. Appropriateness depends not just on the relation between *a* and any possible value that *p* may advance or diminish. It depends, specifically, on what I call *internal values*: values that play a distinctive role in vindicating *p*.[42]

How to mark out these internal values from the set of possible values that blaming might realise? Notice that as an intentional activity, blaming seems to have a basic structure, captured by its minimal description: 'A blames Z for *x*-ing'. That minimal description involves one agent (A, the blamer) responding to a specific moral failing (*x*-ing) in another agent (Z, the blamee). What counts as an internal value is a question

[41] I begin sketching this framework in Upadhyaya 2021, ch. 2.

[42] Compare Tadros' (2020) discussion of 'internal standards of correctness'. Also compare MacIntyre 1981, pp. 187–190.

of what an ideally-motivated version of the person engaging in this activity (A) would pursue within that basic structure.

Here are some candidate answers that I find reasonable. One good the blamer might pursue is that they themselves recognise and appreciate the subject matter to which their blaming activity responds: *A knows the moral facts about Z's x-ing*. Another good the blamer might pursue is that the blamee appreciates those same moral facts: *Z appreciates the morality of their own x-ing*. A third is that the blamee gains moral motivation from that understanding*: Z does what Z ought to do in response to the reasons created by Z's responsibility for x-ing*. Fourth, other people implicated in the wrongdoing, most prominently victims, should be duly recognised as claimants of the knowledge and motivation sought: *where the reasons created by Z's x-ing are directed towards the interests of other agents* (*victims of x-ing*) *both A and Z orient their thinking and behaviour towards those people*. Fifth, the blamer might also pursue the ideal that these facts about wrongdoing are mutually understood through an interactive *accountability process*, in which the parties listen and answer to one another's normative views about the topic at hand: *A*'s *and* Z's *responses to x-ing are shaped through a mutual exercise of A's and Z's normative capacities, as well as those of other agents implicated in A's x-ing*.

The reason the above set of goods appeals to me is that blaming, as I'm minded to conceive of it, has two distinctive vindicating features: *agential* and *procedural*. Blaming is agential in that it is part of aiming to ensure that people respond to moral demands (in blaming, these are demands created by responsibility for wrongdoing) for the right, or morally operative, reasons. Those agential goods set standards for the outcomes blamers ideally seek. For instance, we would hope that Zara does the right thing not because Zara is scared, or minded to avoid penalties, or has been conditioned against recividation; instead Zara should avoid recidivation because of their moral duties to do so. That broad aim – a response to wrongdoing based on morally operative reasons – also sets standards for the processes we might deploy to engage a person through blame. In the best case, a wrongdoer achieves moral understanding and motivation *through* the exercise of their normative capacities. Getting parties to understand these things via other processes, like trickery, pressure, conditioning, or more fantastical means of knowledge-implantation, are non-ideal.

None of the above is a decisive substantive argument for particular internal values. I invite readers to make their own laundry list of blame's distinctive vindicating features. My hope is that having done so, readers will have found the tools for navigating their own response to the problem of appropriate piling-on.

Aside from offering that toolkit, a second advantage of an internal values approach is that it naturally explains why a wide variety of facts seem salient to appropriateness.

We observed several of these in Section II: it seems inappropriate to blame the innocent; to blame people using bad evidence; and to blame for non-moral ends like self-promotion or scapegoating. Those observations make good sense in the light of internal values. When the blamer is unresponsive to the operative moral facts, whether because of ignorance of, or lack of concern for, those facts, they are not intending to develop an accurate understanding of the wrongdoing in question. Nor is it likely an aim of theirs to engage the blamee to reflect on and appreciate those same facts, or to stand up for any of the blamee's victims. These are patent failures to meet the list of internal standards just enumerated – no coincidence, of course, on the Internal Values Account.

A third merit of the account is its explanation of considerations that are *not* salient to appropriateness. Suppose, for example, that blaming a person were to advance goods such as creating happy people; or making people financially better-off; or making a space aesthetically more pleasing. While apt blame may sometimes promote these things, the Internal Values account would deem that fact coincidental. As weighty as those goods might be in deciding how to act, they are not in themselves aptness-affecting since they are not especially connected to blaming. They are, in other words, not internal but *external* values, located outside the basic structure of blaming activity. That is a plausible reason to exclude them.

What's more, the account captures a peculiar, but central pair of explananda: while aptness is a normative category, it must be a rather idiosyncratic one. It seems that we can have moral reasons to act appropriately. But as indicated in Section I, their normative force is distinct from paradigmatic moral reasons – it is clear that reasons not to torture people, for example, or reasons to pay taxes, or reasons to rescue the endangered, differ from reasons to act appropriately. Unlike those reasons, reasons of aptness are undergirded by standards of *felicity*; standards that tell us whether we are performing an activity well (*bona fide*), or not so well (*mala fide*). The reasons are unique to the point of the practice being undertaken.

The Internal Values account readily identifies what that point of a practice is (if it has one): its point is to instantiate the values that play a distinctive role in vindicating it. At the same time, the account explains why that point can be normative. The values that play a role in vindicating a practice are genuine values. They can be the source of moral reasons – roughly, reasons to perform a morally good activity well.

B. Harm Prevents Vindication

Armed with this account, we can begin to see why what the Proportionality View holds might be true: that the harms of piling-on can make it inapt.

Return to the idea that the aptness of *a* within *p* depends on the extent to which *a* advances the values of *p*, and that the values of *p* are internal, i.e. distinctive in vindicating *p* (as set by the basic intentional structure of activity within *p*). Sometimes those values can fail to vindicate *p*. When the values of doing *a* fail to vindicate *p*, the values of *p* are not well-pursued, making *a* inapt.

Harm caused by an activity is among the facts that can prevent a practice from being vindicated in this way. Consider again the four types of disproportionality in blaming – its harms might be gratuitous, counterproductive, unnecessary, and (straightforwardly) disproportionate. Gratuitous blaming involves pointless harm; harm that achieves no good at all, never mind internal goods. Counterproductive blaming does the opposite of contributing to internal value. It actively diminishes the blamee's, and in some instances also the blamers', and victims', abilities to process and redress wrongdoing in an effective way. Unnecessary and disproportionate blaming are further categories of failure to advance internal value. Any internal value achieved through accountability supports harming those called to account only to the extent that is necessary to realise these values. And they only support harms whose disvalue doesn't normatively outweigh their value.

All forms of unjustified harm, then, are ways in which considerations of harm prevent the sound performance of blaming practices, because they prevent internal value from vindicating those practices. That is why those considerations can undermine the aptness of piling-on.

So much for the positive case for the Proportionality View. Section II prefaced a set of concerns about the fact that it relates the aptness of blaming to its outcomes. We are now better poised to take those up.

V. APTNESS & OUTCOMES

The aptness of blaming obviously comes apart from its good and bad effects. We raised the examples of inapt blaming to excellent reformative effect (piling-on to make an example of one wrongdoer in order that many others reform), and apt blaming with little to no reformative effect (blaming psychopaths or the morally hardened). How does the argument given handle these problems?

First recall the idea that blame's internal value is limited by the activity's minimal description – 'A blames Z for *x*-ing'. Apt blaming is thus not a matter of just any good effect: the goods to be effected within that structure must relate closely to A, Z, and those closely connected to Z's *x*-ing. Although reform of other wrongdoers is a good outcome of blaming, that outcome does not relate closely to the blamee's understanding of their own reformative duties. To put it differently, reform of people agentially distant from

Z's x-ing is another instance of an external value of blaming, located outside of the activity's basic structure. As such, it does not contribute to aptness on the Internal Values Account.

Second, the account allows that aptness is *multidimensional*. Because blaming has more than one internal value, the practice might be defective with respect to one such value while fulfilling another. Thus, while blaming the morally hardened might be defective in securing their reform, it might nonetheless instantiate other values like the blamer's understanding of their wrongdoing, and acknowledgment of the claims of victims who belong to the history of that wrongdoing. Blaming the morally hardened can be apt in these latter respects despite being inapt in the former.

That second element of the account sets it apart from *fittingness accounts*. Fittingness accounts hold that aptness is a 'relation of fit' between a particular response, on the one hand, and the facts to which it responds, on the other. For example, it is a fact about Mount Everest that it is beautiful. Responses which 'fit' that fact – e.g. smiling with wonderment at the way the sunlight reflects off the crevices in the rockfaces – are apt. Responses that do not 'fit' – e.g. recoiling with disgust – are inapt. Or consider a funny joke. Sobbing inconsolably is inapt because it does not 'fit'; laughing is apt because it 'fits'. Whether something is apt, on this view, is just a question of whether it constitutes an act that 'fits', rather than about outcomes. Applied to blame, suppose one holds that some response, e.g. disappointment or anger, 'fits' a responsible act of wrongdoing. In that case, whether blaming is apt is just the question of whether it *is* a kind of disappointment, anger, etc. The good or bad outcomes of that response are neither here nor there. This in part is why Amia Srinivasan holds that counterproductive anger can be apt.[43] Srinivasan argues that appreciation of wrongdoing is good in itself, regardless of what this does to redress wrongdoing.

Some might have worried that relating aptness to harmful outcomes comes at the cost of denying that insight – that certain reactions are apt just because they constitute appreciation of moral violation. However, the Internal Values Account concurs that appreciation of wrongdoing is an internal value. It only adds that aptness is multidimensional. A blaming practice can be apt in instantiating each person's appreciation of wrongdoing. It can, at the same time, be inapt when it comes to instantiating other internal values, for example the blamee's shared appreciation of that wrongdoing, and their redressive motivation based on it.

[43] Srinivasan 2017. See also D'Arms & Jacobson 2003 (though at p. 143 D'Arms & Jacobson do acknowledge the importance of other factors). And compare Owens 2012, ch. 2.

This point helps to substantiate my contention in Section III that beyond adopting fitting attitudes, it can also be apt to seek felicitous communication of those attitudes.[44] There is indeed reason to view both values as complementary. Appreciation of wrongdoing is after all incomplete without appreciation of the duties of reform that wrongdoing creates.[45] Something would be amiss if a *blamee*, upon gaining that appreciation, lacked any motivation to make amends. But, for the same reasons, it also wouldn't be ideal if a *blamer* appreciated an instance of wrongdoing with nil motivation to ensure that this understanding is mutually recognised and acted upon, by urging the blamee to apply their agential capacities to its subject matter. To be sure, whether communication is ultimately felicitous in gaining that uptake is a question about the outcomes of accountability interactions. But these outcomes derive some of their value from the fact that they are brought about by a joint agential procedure aimed at mutual understanding.

In the light of this discussion, let's return to the general challenge: What is the difference between blaming's being apt and its simply being a good idea in view of its expected outcomes? Whether blaming is apt depends *only* on the standards for realising internal values. As noted, some of those values relate to outcomes, but others are procedural. In contrast, whether blaming is a good idea depends on its *full range* of expected outcomes. The full range can involve internal values (at least those that are outcome-related) but also external values, as well as any other normative feature of outcomes.

It follows that doing something apt can sometimes be a good idea – when and because instantiating internal values happens also to be recommended by its full range of expected outcomes. But that co-occurrence is never guaranteed. It could always be a good idea to do something inapt – because the full range of outcomes could always recommend doing something other than instantiating internal values.

VI. INTERNAL & EXTERNAL PROPORTIONALITY

Having worked through an account of the relationship between the aptness of blaming and disproportionate outcomes, we might be left wondering how it differs from prominent accounts of proportionality in online criticism.[46]

---

[44] I thank an anonymous referee for helping me clarify this point. For slightly different variations on this same theme, see footnote 8, as well as Mason 2019, ch. 5, and Macnamara 2015.

[45] Sliwa (2020) refers to this more broadly as wrongdoing's 'normative footprint'.

[46] An anonymous referee prompted this discussion.

Billingham and Parr in particular discuss proportionality in detail. They focus on justification conditions of public shaming, arguing that shaming is unjustified if it is (inter alia) counterproductive, or unnecessary, or causes harms outweighing its goods. The Proportionality View holds that the same objections to a blaming practice can make it inapt.

The two accounts appear alike if we subscribe to the idea that blaming and shaming can do similar harm. I suggested in Section III.A that blaming can do harm even if each blamer intends not to shame but sincerely to call someone to account. I suggested in Section III.C that blaming can do harm even if each blamer does not communicate with the wrongdoer. There's no reason to exclude these harms from the Billingham-Parr proportionality conditions of justification. It might seem, therefore, that all blaming practices condemned by those conditions are inapt, and vice-versa, that all blaming practices condemned by the Proportionality View are also unjustifiable.

But the two views in fact differ extensionally. Whereas the Proportionality View is indexed to a subset of values, the Billingham-Parr view is not. The Proportionality View constrains the pursuit of internal values only. It sets, if you like, a standard of *internal proportionality* which determines whether a practice's ills are morally supported from the inside; or supported by the point or purpose of the practice. The Billingham-Parr view constrains the pursuit of the full range of possible outcomes of public criticism. That, I noted in the previous section, is a much wider-scope constraint, involving both internal and external values, as well as other normative facts affecting the overall justification of criticism. We can call it a standard of *external proportionality*: a standard that determines whether we ought to tolerate a practice from the outside, whether or not it is morally supported from the inside.

Billingham and Parr carve out a narrower standard that might again appear closer to internal proportionality. They distinguish between *narrowly* and *widely proportionate* criticism. Narrowly proportionate criticism harms people less than or as much as they are *liable* to be harmed – its harms needn't infringe anyone's rights. Widely proportionate criticism harms people beyond the extent to which they are liable: while that criticism infringes people's rights, it is proportionate in the sense that the goods achieved by infringing those rights are important enough to justify doing so.[47]

But this distinction between narrow and wide proportionality in blaming also comes apart from the distinction between internal and external proportionality. Were they the same distinction, all (and only all) harm that meets proportionality standards

[47] I intend this as a suitably neutral statement of the distinction drawn by McMahan (2009, pp. 20–21).

for appropriate blaming is harm that a blamee is liable to suffer: harm that does not infringe the blamee's rights. Yet that is not true, at least according to the argument of this essay.

Granted – much harm that we do to a person through appropriately blaming them is harm that the blamee is liable to suffer. That much is true for the simple reason that wrongdoers incur duties to be accountable for their wrongdoing in virtue of having committed it. These are the duties I've been emphasising as part of internal value – to reflect on what one did, and to be motivated by that reflection to meet additional duties like apology, compensation, and reform. A wrongdoer's rights are not normally infringed by the costs of being blamed in the service of those duties.

However, responsibility for wrongdoing can also make a person liable to be harmed for non-blaming reasons. Suppose, for example, that a person crushed by a dogpile has been harmed to a much greater extent than is necessary for them to achieve moral understanding and redressive motivation, and for a community to speak out for the victim. Suppose, though, that this already overblamed person has yet to compensate their victim. Coercive sanctions might then be narrowly proportionate despite being internally disproportionate. Forcibly exacting compensation from the wrongdoer might not infringe their rights, for example, yet this fact would not undo the proportionality defects within the blaming practice itself.

Last, suppose that a person is overpunished for their wrongdoing: they have, let's say, served most of a wildly excessive prison sentence. But suppose that until now, this person has never been called to account by the victim of their wrongdoing. It seems to me that this victim could be vindicated by internal value in arranging a visit to exhort the wrongdoer to apologise and explain themselves, even if that process heaps more costs upon the already overpunished blamee. So, it seems, a person might be blamed a way that satisfies internal proportionality, but not narrow proportionality. Though the overpunished blamee's rights not to be harmed have long been infringed, the blaming practice remains felicitous as the proportionality defects are located outside of it.

In sum, proportionality in apt blaming determines how well-performing our blaming practice is; whether it is a defective, or morally corrupted, version of itself. That is how its implications differ from conceptions of proportionality often applied to content moderation and online criticism. Those conceptions are about the wider justification of these practices, independently of whether they have been corrupted in some way.

Some might wonder – how significant is the internal category we have just carved out?[48] External proportionality calculations, which count both internal and external

[48] Thanks to the Editor for discussion.

values, seem likely to dominate internal proportionality calculations in determining how to blame people together. Does this mean that internal proportionality is normatively trivial or inert? Not necessarily. After all, a good deal is morally at stake within our accountability practices. Wrongdoers have powerful obligations to be answerable to others, and to deliver certain responses. The wider community has powerful obligations to acknowledge the facts about wrongdoing and the interest of victims. Given the importance of those things, we have at least some presumptive moral reason to honour them: to strive for their felicitous pursuit. No doubt external considerations will often compel us to do otherwise. But at a cost.

VI. CONCLUSION

Time to take stock of the main moves of this essay, by reminding ourselves of the problem it began with:

i. It is appropriate for each person (identically situated) among many to blame a person who is responsible for wrongdoing.

ii. If it is appropriate for each person (identically situated) among many to blame a wrongdoer, then it is appropriate for everyone to blame that wrongdoer, regardless of how many blamers there are.

i and *ii* preclude:

iii. It is not appropriate for everyone (identically situated) to blame a wrongdoer if the number of blamers is sufficiently high.

But *iii* is plausible too.

I have argued that *ii* is the least plausible, because apt blaming is subject to internal proportionality standards: standards which multiple agents together can flout. As we can make sense of harm aggregating across many people's acts, internal values can lose vindicating force through harm that no single piler-on inflicts on their own. The argument thus leaves room for the possibility both that each piler-on's contribution is apt, and that everyone piling-on is inapt.

Albeit a widget for closing that particular door, I fear that last possibility opens a new one. When it obtains – when it is true both that each blames well, and that everyone blames badly – one dimension of internal value is secured at the expense of another.[49]

[49] I thank an anonymous referee for this observation.

Each piler-on's sound blaming performance undermines everyone's sound blaming performance. This poses a problem about whether to prioritise individual or collective conduct. On the one hand, all of the pilers-on morally corrupt the practice. Yet on the other, no piler-on's behaviour is corrupted, and accountability goods are not well-advanced if none of them blames. Is it most apt, in these situations, that everybody piles-on, or that nobody does? I submit this question for future inquiry, in the vague hope that collective action theory might help resolve it.

ACKNOWLEDGMENTS

Thank you for brilliant comments to Chris Mills, Massimo Renzo, Victor Tadros, two anonymous referees and the Editor of this journal, as well as seminar discussants at the Department of Applied Philosophy & Ethics, Institute of Philosophy, Academy of Sciences, Prague. This work was supported by the National Research Centre on Privacy, Harm and Adversarial Influence Online (REPHRAIN); and the Lumina Quaeruntur Grant of the Czech Academy of Sciences (LQ300092001).

COMPETING INTERESTS

The author declares that there are no competing interests.

REFERENCES

Bennett, Christopher. 2002. The varieties of retributive experience. *Philosophical Quarterly*, 52: 145–163. https://doi.org/10.1111/1467-9213.00259

Bhargava, Vikram. 2020. Firm responses to mass outrage: technology, blame, and employment. *Journal of Business Ethics,* 163: 379–400. https://doi.org/10.1007/s10551-018-4043-7

Billingham, Paul and Tom Parr. 2020. Enforcing social norms: the morality of public shaming. *European Journal of Philosophy*, 28: 997–1016. https://doi.org/10.1111/ejop.12543

D'Arms, Justin and Daniel Jacobson. 2003. The significance of recalcitrant emotion. *Royal Institute of Philosophy Supplement*, 52: 127–145. https://doi.org/10.1017/S1358246100007931

Driver, Julia. 2015. Individual consumption and moral complicity. Pp. 67–79 in *The Moral Complexities of Eating Meat*, ed. Ben Bramble and Bob Fischer. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199353903.003.0005

Duff, Antony. 2011. Blame, moral standing and the legitimacy of the criminal trial. *Ratio*, 23: 123–40. https://doi.org/10.1111/j.1467-9329.2010.00456.x

Ellen, Barba. 2023. Once shamed, forever a s***. *The Guardian Online*, 8 Oct. https://www.theguardian.com/commentisfree/2023/oct/08/once-shamed-forever-a-slut-doomed-to-stay-in-the-internet-spotlight

Fricker, Miranda. 2016. What's the point of blame? A paradigm based explanation. *Noûs*, 50: 165–183. https://doi.org/10.1111/nous.12067

Fritz, James. 2021. Online shaming and the ethics of public disapproval. *Journal of Applied Philosophy*, 38: 686–701. https://doi.org/10.1111/japp.12510

Frye, Harrison. 2021. The problem of public shaming. *Journal of Political Philosophy*, 30: 188–206. https://doi.org/10.1111/jopp.12252

Howard, Jeffrey. 2019. Dangerous speech. *Philosophy & Public Affairs*, 47: 208–254. https://doi.org/10.1111/papa.12145

Howard, Jeffrey. Forthcoming. *Setting Fire to Reason: The Ethics of Free Speech.* Princeton, NJ: Princeton University Press.

Kagan, Shelly. 2011. Do I make a difference? *Philosophy and Public Affairs*, 39: 105–14. https://doi.org/10.1111/j.1088-4963.2011.01203.x

Kelly, Erin. 2018. *The Limits of Blame*. Cambridge, MA: Harvard University Press https://doi.org/10.4159/9780674989436

Kutz, Christopher. 2000. *Complicity: Ethics and Law for a Collective Age*. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511663758

MacIntyre, Alistair. 1981. *After Virtue*. Notre Dame, IN: University of Notre Dame Press. https://undpress.nd.edu/9780268035044/after-virtue/

Macnamara, Coleen. 2015. Reactive attitudes as communicative entities. *Philosophy & Phenomenological Research*, 90: 546–569. https://doi.org/10.1111/phpr.12075

Mason, Elinor. 2019. *Ways to be Blameworthy: Rightness, Wrongness & Responsibility*. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780198833604.001.0001

Max, D. T. 2020. The public shaming pandemic. *New Yorker*, 28 Sept. https://www.newyorker.com/magazine/2020/09/28/the-public-shaming-pandemic

McGeer, Victoria. 2019. Scaffolding agency: a proleptic account of the reactive attitudes. *European Journal of Philosophy*, 27: 301–23. https://onlinelibrary.wiley.com/doi/10.1111/ejop.12408

McKenna, Michael. 2012a. *Conversation and Responsibility*. New York: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199740031.001.0001

McKenna, Michael. 2012b. Directed blame and conversation. Pp. 119–140 in *Blame: Its Nature and Norms*, ed. D. Justin Coates and Neal A. Tognazzini. New York: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199860821.003.0007

McMahan, Jeff. 2009. *Killing in War*. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199548668.001.0001

McMahan, Jeff. 2017. Liability, proportionality, and the number of aggressors. Pp. 3–27 in *The Ethics of War: Essays*, ed. Saba Barzagan and Samuel Rickless. New York: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199376148.003.0001

Nefsky, Julia. 2023. Participation, collective impact, and your instrumental significance. *Journal of Practical Ethics*, 11: 38–58. https://doi.org/10.3998/jpe.4624

Nelkin, Dana. 2016. Accountability and desert. *Journal of Ethics*, 20: 173–189. https://doi.org/10.1007/s10892-016-9230-0
Norlock, Kathryn. 2017. Online shaming. *Social Philosophy Today*, 33: 187–197. https://doi.org/10.5840/socphiltoday201762343
Owens, David. 2012. *Shaping the Normative Landscape*. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199691500.001.0001
Parfit, Derek. 1984. *Reasons & Persons*. Oxford: Oxford University Press. https://doi.org/10.1093/019824908X.001.0001
Pereboom, Derk. 2014. *Free Will, Agency, and Meaning in Life*. New York: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199685516.001.0001
Pickard, Hannah. 2011. Responsibility without blame: empathy and the effective treatment of personality disorder. *Philosophy, Psychiatry, Psychology*, 18: 209–224. https://doi.org/10.1353/ppp.2011.0032
Pickard, Hannah. 2013. Irrational blame. *Analysis*, 73: 613–626. https://doi.org/10.1093/analys/ant075
Pressly, Lowly. 2024. The right to be forgotten and the value of an open future. *Ethics*, 135: 65–87. https://doi.org/10.1086/731431
Radzik, Linda. 2011. Minding your own business: differentiating accountability relations within the moral community. *Social Theory & Practice*, 27: 574–98. https://doi.org/10.5840/soctheorpract201137434
Rossi, B. 2018. The commitment account of hypocrisy. *Ethical Theory & Moral Practice*, 21: 553–67. https://doi.otg/10.1007/s10677-018-9917-3
Seim, Maria. 2019. The standing to blame and meddling. *Teorema*, 38: 7–26. https://philarchive.org/archive/SEITST-2v1
Shoemaker, David and Manuel Vargas. 2019. Moral torch-fishing: a signalling theory of blame. *Noûs*, 55: 581– 602. https://doi.org/10.1111/nous.12316
Slavny, Adam. 2017. The normative foundations of defamatory meaning. *Law & Philosophy*, 37: 523–547. https://doi.org/10.1007/s10982-017-9321-5
Sliwa, Paulina. 2020. Reverse-engineering blame. *Philosophical Perspectives*, 33: 200–219. https://doi.org/10.1111/phpe.12131
Srinivasan, Amia. 2017. The aptness of anger. *Journal of Political Philosophy*, 26: 123–144. https://doi.org/10.1111/jopp.12130
Strawson, Peter. 1962. Freedom and resentment. *Proceedings of the British Academy*, 48: 187– 211. https://doi.org/10.4324/9780203882566
Tadros, Victor. 2020. Criminalisation: in and out. *Criminal Law & Philosophy*, 14: 365–280. https://doi.org/10.1007/s11572-020-09536-7
Thomason, Krista. 2021. The moral risks of online shaming. Pp. 145–162 in *The Oxford Handbook of Digital Ethics*, ed. Carissa Véliz. Oxford: Oxford University Press. https://doi.org/10.1093/oxfordhb/9780198857815.013.8
Todd, Patrick. 2017. A unified account of the moral standing to blame. *Noûs*, 53: 347–374. https://doi.org/10.1111/nous.12215
Tomlin, Patrick. Forthcoming. *Violence in Proportion*. Oxford University Press.

Tosi, Justin and Brandon Warmke. 2020. *Grandstanding: The Use and Abuse of Moral Talk*. New York: Oxford University Press. https://doi.org/10.1093/oso/9780190900151.001.0001

Upadhyaya, Kartik. 2021. What's wrong with hypocrisy. PhD Thesis. Department of Philosophy, University of Warwick. https://wrap.warwick.ac.uk/id/eprint/160547/

Upadhyaya, Kartik. 2024. Standing, more or less...Unpublished paper, University of Warwick.

Wallace, R. Jay. 1994. *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press. https://doi.org/10.2307/2956371

Wallace, R. Jay. 2019. Moral address: what it is, why it matters. Vol. 5, pp. 88–109 in *Oxford Studies in Agency and Responsibility*, ed. Justin D. Coates and Neal A. Tognazzini. Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780198830238.003.0005

Wieland, Jan Willem. 2022. Participation and degrees. *Utilitas*, 34: 39–56. https://doi.org/10.1017/S095382082100025X