

Understanding Degrees of Responsibility

1. Preliminary Remarks

It is something of a truism that agents can be more (or less) morally responsible for their actions. By “truism” I don’t mean, of course, “something that everyone accepts.” Rather, I simply mean that our practices of holding others morally responsible seem to have, as an underlying assumption, a commitment to the idea that agents can be more (or less) praiseworthy or blameworthy for their actions. And though philosophers are famously given to trying to upset just this sort of underlying assumption, I will not be doing that here. For I think that it is true: agents *can* be more (or less) praise- or blameworthy for their actions. What I seek to do, instead, is to understand what precisely this assumption really comes to. As we’ll see, the assumption is actually underwritten by several independent properties in virtue of which an agent might be more (or less) responsible—a fact that has important downstream implications.

Discussions that concern degrees of responsibility, praise- and blameworthiness can get unwieldy pretty quickly. So as a way of simplifying, I want to focus exclusively on blameworthiness. In particular, I want to make sense of degrees of blameworthiness—i.e., what it is for an agent to be more (or less) blameworthy for committing some transgression. I restrict myself in this way for two reasons. First, it is easier and more parsimonious to simply talk about “blameworthiness” than it is to talk about “moral responsibility and praise- and blameworthiness” (or even “responsibility, et al.”). Of course, it seems plausible that what I say about blameworthiness might go for moral responsibility more generally. So if one is so inclined, one can simply substitute talk of “blameworthiness” for “moral responsibility for a transgression.” If one doubts that such a substitution can be made, then

fair enough. But even in the event that no substitution can be made, we'll still have an account of what it is to be more (or less) blameworthy.

Second, it's not altogether clear that praise- and blameworthiness are different only in their valences. Indeed, they seem to be *asymmetrical* in a number of ways. For instance, some have suggested that while praiseworthiness doesn't require the ability to do otherwise, blameworthiness does.¹ A second way that praise- and blameworthiness might be asymmetrical is due to the fact that the phenomenon of praise and blame themselves might be quite different. Whereas blame can be wholly internal—I can blame you “in my heart,” even if I never communicate that blame—it seems strange to think of wholly internal forms of praise. Gratitude, to be sure, can be uncommunicated. But while resentment seems to be at the heart of what it is to blame, gratitude seems less closely related to praise. I might, after all, feel grateful to you even if I don't think you're praiseworthy for your action. (As in the case in which you do something to help me but for purely self-interested reasons.) So too, I can praise you without experiencing gratitude. (As in the case in which I laud and commend your actions to others, but because idiosyncratic interests of my own, I also wish that you hadn't done those things.) More simply: it just seems strange to say that I've praised you for your action if I've never communicated that. This doesn't seem to be the case for blame, however.

A third potential asymmetry between praiseworthiness and blameworthiness is that the activities of praise and blame are apparently governed by asymmetrical norms. Most notably, blame seems apparently governed by norms of fairness in a way that praise isn't. It

¹ Cf. Susan Wolf, “Asymmetrical Freedom,” *The Journal of Philosophy* 77.3 (1980): 151-66; Susan Wolf, *Freedom Within Reason*, New York: Oxford University Press (1990); and Dana K. Nelkin, *Making Sense of Freedom and Responsibility*, New York: Oxford University Press (2011). For reasons to reject this purported asymmetry, see John Martin Fischer and Mark Ravizza, “Review: Responsibility, Freedom, and Reason,” *Ethics* 102.2 (1992): 368-389.

is possible, of course, to praise someone unfairly, as in the case in which you praise Sarah for work that Mitchell did. But it's less clear that another kind of unfairness arises in the case of praise. For wholly noncomparative reasons, it sometimes seems unfair to blame an agent for some action, even though it wouldn't seem unfair to praise that agent if one were to hold fixed everything about the situation except the valence of her action.

For these reasons, I'm not sure that an account of what it is to be more (or less) blameworthy is isomorphic to an account of what it is to be more (or less) praiseworthy. (I suspect that it might be, but I'm just not prepared to settle that matter here.) Accordingly, I plan to simply avoid these delicate issues and instead focus exclusively on blameworthiness. If one isn't too worried about the alleged asymmetries between praise- and blameworthiness, then please feel free to extend the account to praiseworthiness. However, if one is worried about asymmetries then again, you'll at least get a theory of what it is to be more or less blameworthy for an action. And that ain't nothing.

Finally, I want to specify exactly what I will (and what I won't) do in this paper. In giving an account of what it is to be more (or less) blameworthy, I won't be giving an account of the conditions under which we are more (or less) blameworthy for our actions. Recently, a number of philosophers have given us theories about the conditions under which agents are more (or less) blameworthy.² I find much to recommend in these accounts, and so nothing I say here should be seen as an attack on them. Instead, what I'm doing in this paper is something more fundamental. Whereas these other philosophers have been interested in telling us the conditions under which an agent instantiates the property of being

² Cf. D. Justin Coates and Philip Swenson, "Reasons-responsiveness and Degrees of Responsibility," *Philosophical Studies* 165.2 (2013): 629-645; Manuel Vargas, *Building Better Beings*, New York: Oxford University Press (2013); and Dana K. Nelkin, "Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness," *Nous* (forthcoming). For some worries about these practices, see Jesper Ryberg, "Punishing Adolescents – On Immaturity and Diminished Responsibility," *Neuroethics* 7.3 (2014): 327-336; Jesper Ryberg, "Responsibility and Capacities: A note on the Proportionality Assumption," *Analysis* 74.3 (2014): 393-397.

more (or less) blameworthy, I'm interested in telling you what exactly that property (or properties) comes to. To be sure, how we answer the latter question will have implications for how we answer the former. But I do not pursue that issue in detail in this paper. Here, I simply want to offer a theory of degrees of blameworthiness that both makes sense from a normative point of view and also can underwrite the truth of the truism I started with.

2. Blameworthiness

If we're going to crack the nut of degrees of blameworthiness, it makes sense to start with a statement of what it is to be blameworthy. Blameworthiness, as I understand it, has a descriptive and normative element. If one is blameworthy, then one is *worthy* of something (that's the normative bit of blameworthiness), namely *blame* (this is the descriptive bit). But that doesn't yet tell us what it is to stand in the *being worthy of* relationship. Nor does it tell us what it is to blame.

Concerning the former issue, blameworthiness is most commonly cashed out in terms of being *deserving* of blame. But what does this mean? Does it simply mean, as some have suggested, that blameworthiness is a matter of blame being an appropriate to one's action(s)?³ I think not. For though there are a number of similarities between desert and other normative relationships like appropriateness, the notion of desert at stake is distinct. After all, it might be inappropriate to blame someone, even if they deserve it.⁴ This is true, for example, in cases of complicity. If *I'm* the one who convinces you to break your promise, it's not appropriate for me to blame you for doing so. But this doesn't mean that you don't *deserve* to be blamed. So desert and appropriateness are different, and while it's important that

³ Cf. R. Jay Wallace, *Responsibility and the Moral Sentiments*, Harvard University Press (1994).

⁴ Cf. Angela Smith, "On Being Responsible and Holding Responsible," *Journal of Ethics* 11.4 (2007): 465-484.

our blame is appropriate, blameworthiness concerns the more circumscribed relationship of *deserving* blame.

A second thing to note concerning the nature of desert is that standing in the desert relationship is not reducible to consequentialist or contractualist reasons. I am not deserving of blame, the thought goes, simply because it would be *good* to blame me. Nor am I deserving of blame simply because I cannot reasonably complain about being targeted with blame. To be sure, if it is good to blame me, then it might be that one should, all things considered, do so. So too, if I cannot reasonably reject the principle that allows you to blame me, then again, blame might be all things considered the right thing to do. But it is consistent with the possibility that in neither of these scenarios do I deserve to be blamed in some more basic sense.⁵ And it is this more basic sense of desert that matters for *blameworthiness*. Being more (or less) blameworthy is thus a matter of being more (or less) deserving for blame.

This tells us about the normative bit of blameworthiness, but by itself, it doesn't give us the full story of what it is to be blameworthy. To see this, let's compare two of the many extant theories of blame. On the first theory, A blames B for her transgression by judging that B has acted in a faulty manner. On the second theory, A blames B for her transgression by resenting B (or by experiencing some other form of anger towards B) for her transgression. The property of being deserving of a judgment that one has acted in a faulty manner is conceptually distinct from the property of being deserving of resentment or other forms of anger. As a result, these differing accounts of what it is to blame give rise to

⁵ For a clearer statement of "basic desert" or "desert in the basic sense," see Derk Pereboom, *Free Will, Agency, and Meaning in Life*, New York: Oxford University Press (2014); see also Michael McKenna's discussion of these issues in McKenna, *Conversation and Responsibility*, New York: Oxford University Press (2012). For the remainder of the paper, when I speak of "desert" I mean to refer exclusively to desert in the basic sense.

accounts of blameworthiness are also different. To understand what it is to be blameworthy, it seems that we'll first have to understand what it is to blame.

Unfortunately, understanding what it is to blame is itself a quite difficult task. In the past two decades, no less than R. Jay Wallace, George Sher, T. M. Scanlon, Susan Wolf, Stephen Darwall, Angela Smith, and Michael McKenna (to mention only a few) have weighed in on this issue without coming to much consensus.⁶ If these top rate folks can't settle the issue over the course of hundreds (if not thousands) of pages, I can't reasonably hope to do so in a couple passing remarks.⁷ So rather than trying to wade into the issue of what blame comes to, I'm going to simply avoid it. No doubt, precisely identifying the conditions under which an individual instantiates the property of being more (or less) blameworthy require that we answer the question of what it is to blame. But this issue is less pressing when we're just trying to understand what it is to be more (or less) blameworthy. This is because understanding the gradations in the desert relationship (and not the nature of blame itself) is really what matters for the purposes of understanding what it is to be more (or less) blameworthy for an action. As I see it, then, the theory of what it is to be more (or less) blameworthy that I'll offer in the remainder of the paper is "plug and play" with respect to a theory of blame. That is, take your preferred theory of blame, stick it into the analysis, and voila! You have your own custom theory of what it is to be blameworthy.

⁶ Wallace (1994); George Sher, *In Praise of Blame*, New York: Oxford University Press (2006); T. M. Scanlon, *Moral Dimensions: Permissibility, Meaning, Blame*, Cambridge: Harvard University Press (2008); Susan Wolf, "Blame, Italian Style," in *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*. eds. R. Jay Wallace, Rahul Kumar, and Samuel Freeman, New York: Oxford University Press (2011); ; Stephen Darwall, *The Second-Person Standpoint: Morality, Respect, and Accountability*, Cambridge: Harvard University Press (2006); Angela Smith, "Moral Blame and Moral Protest," in *Blame: Its Nature and Norms*, eds. D. Justin Coates and Neal A. Tognazzini, New York: Oxford University Press (2013): 27-48; McKenna (2012).

⁷ For what it's worth, my own view, owing to the fact that "blame" has such a wide semantic range, is that blame is actually a *genus* of normatively significant responses and that almost all theories of blame genuinely identify distinct *species* of blaming responses.

Being blameworthy is thus being deserving of blame, whatever that comes to. Consequently, I will now talk not of blameworthiness but of being deserving of blame. The issue, then, is ultimately a matter what it means to be more (or less) deserving of blame. At first blush, this might seem straightforward. But as we'll see, the property (or properties) in question doesn't (or don't) wear its content on its sleeve.

3. Ambiguities

In common practice, talk of degrees of blameworthiness or of being more (or less) blameworthy is loose. Here's one clear way in which it's loose. How does one express that an adult criminal offender is more blameworthy than he would have been if he had committed a type-identical crime at the age of thirteen? Well, since blameworthiness is fundamentally a matter of desert, you might naturally think that the adult criminal offender is more deserving of blame (and punishment) than he would have been if he had committed a type-identical crime at the age of thirteen. This sounds nice enough. But the notion of "more deserving" is multiply ambiguous.

On the one, perhaps the most natural, reading of "more deserving," it seems that one is more deserving of some response just in case it is (intrinsically) better, solely as a matter of desert, that that person be responded to in the deserved way. Or, if one is skeptical about talk of intrinsic goodness, you might say that for A to be more deserving of blame for a token of an A-type transgression than B is for a token of an A-type transgression, there are weightier reasons for blaming A than for blaming B, and that those reasons are due wholly to the facts in virtue of which A and B deserved to be blamed.

I say that this is perhaps the most natural reading of "more deserving" because the "more" here seems to point to how strongly the individual stands in the desert relationship.

In other words, the property of being more (or less) deserving of blame seems tied to how weighty the desert-based reasons for blaming are.

And yet, despite the naturalness of this interpretation, I'm doubtful that this is really what most of us are talking about when we talk about A's being more deserving of blame than B.⁸ Weaker: I doubt that this is the only property that most of us are talking about when we talk about A's being more deserving of blame than B. This is clear in the case of adult versus juvenile criminal offenders. I think that many of us, when asked it means to say that an adult criminal offender is more deserving of blame (and punishment) than he would be had he committed a type-identical crime at the age of thirteen, would say that while the adult offender is deserving of m years of imprisonment, the teenager would only be deserving $m-n$ (where n is a positive integer smaller than m) years of punishment. Here, the idea seems to be not that the reasons for blaming (or punishing) the teenager are less weighty, or that it is less intrinsically good to punish the teenager, but that the amount of blame (or punishment) that is deserved is of a greater quantity. And though we often express this idea by saying that the adult offender is more deserving of blame (or punishment) than the criminal offender, it seems more precise to say that the adult is deserving of more blame (or punishment) than the teenager.

This means that there are at least three distinct properties that we could be picking out when we describe someone as being more (or less) blameworthy. We could be claiming (i) that the agent is *more deserving* of blame. Or (ii) we could be claiming that the agent is *deserving of more* blame. Finally, (iii) we could be claiming both—that the agent is more deserving of blame and that he is deserving of more blame than he otherwise would've been. Giving a full account of being more (or less) blameworthy will thus require that we

⁸ N.B. This doesn't mean that I think this property is unimportant!

understand all three of these properties. In the next four sections, then, I want to discuss each property in turn. After that, I'll explore what this multiplication of properties means for our practices of holding people more (or less) blameworthy.

4. Being More Deserving of Blame

As I said above, when I say that A is *more deserving* of blame than B, what I'm claiming is that it is better from the point of view of desert that we blame A than it is that we blame B. Or alternatively: that the desert-based reasons to blame A are weightier than the desert-based reasons to blame B. But what do these proposals really mean?

Recently, Shelly Kagan has developed a visual apparatus for understanding these sorts of proposals: so-called “desert graphs.”⁹ These graphs purport to perspicuously represent what precisely an agent deserves and how much they deserve it (i.e., how good it is from the point of view of desert that they receive whatever it is that they deserve). Furthermore, these graphs are important because, as Kagan says, “there is indeed far more complexity in our ideas about desert than we have yet unearthed,” (Kagan, 2012: 46). Kagan is surely right about this, and representing these issues of Kagan-style graphs enables us to explore the complexity at issue. Indeed, the ambiguities that I introduced in the previous section, and that I will explore in more detail in subsequent sections, were not apparent to me before I thought carefully about the geometry of desert. This suggests, that in the context of discovery, at least, the apparatus of desert graphs have their place. But this undersells their importance, I think. As I said above, in considering these graphs, it is plausible that we can better understand the structure of desert itself.

⁹ Shelly Kagan, *The Geometry of Desert*, New York: Oxford University Press (2012).

To make sense of desert graphs, we must first understand what exactly they represent. This means that we need an interpretation of their axes. Here, I will roughly follow Kagan. However, because he is less interested in *deserved blame* in particular, I will take some liberties in interpreting the graphs' axes. What follows, then, is in the spirit of Kagan, but not entirely faithful to the program he sets forth in *The Geometry of Desert*.

As I will interpret it, the x-axis of Kagan-style desert graphs represents the praising or blaming responses that agents deserve in light of their actions. For Kagan, “the X-axis ... represent[s] whatever it is that people deserve more or less of by virtue of being morally better or worse individuals,” (Kagan 2012: 47). For us, however, the x-axis represents whatever people deserve more or less of by virtue of being more or less morally responsible for their actions.¹⁰ Consequently, points to the right of the y-axis represent positive responses, including praise and other positive contributors to wellbeing. Points to the left of the y-axis represent negative responses, including blame, punishment, and other forms of sanctions. The farther away from the y-axis one goes, the more of any particular response is warranted. This means that in the case of blame, the amount of blame that is deserved will increase as the corresponding x coordinate gets smaller. On this interpretation, then, an x coordinate of -4 represents less blame as being deserved than an x coordinate of -13.


The y-axis of these graphs represents how intrinsically *good it is from the point of view of desert* that an agent receives the response in question (this is Kagan's preferred locution). Or alternatively, they represent how weighty the desert-based reasons for response are (this is my preferred locution—for what it's worth, I think little hangs on the difference between

¹⁰ Since even generally good individuals can sometimes be morally responsible for acting wrongly, there will be important differences between an X-axis that concerns the moral goodness of individuals and one that concerns their moral responsibility.

Kagan's locution and mine). This means that if the blame that an agent deserves is represented by a point that is very high on the y-axis, then it follows that it is very good from the point of view of desert that we blame that individual or that we have very weighty desert-based reasons for blaming that individual. If the potential response is below the x-axis, then it follows that it is bad from the point of view of desert that the agent be blamed or that we have weighty desert-based reasons to refrain from blaming that individual. In other words, points to the left of the y-axis and below the x-axis represent agents who do not deserve blame at all. Since I am concerned exclusively with deserved blame in this paper, however, I'll be focusing on the top left quadrant of these graphs, since this is the area where agents can be represented as being deserving of blame.

With this in mind, let's consider Figure 1, which represents two points, A and B, above the x-axis and to the left of the y-axis.

Fig 1.



In Fig. 1, individual A and B each deserved to be blamed. We can see this because what it is that they deserve is represented as being to the left of the y-axis and above the x-axis (i.e., in the top left quadrant of the Cartesian coordinates). In each case it is good from the point of view of desert that we blame the individual. So too, in each case, we have reason to blame the individual. But as we see here, it is intrinsically better that A be blamed than it is that B be blamed; i.e., there are weightier reasons to blame A than there are to blame B. This graph, then, is a visual representation of the state of affairs we are describing when we attribute to an agent the property of being *more deserving* of blame. And though this is apparently a comparative judgment of how much blame two distinct agents deserve, it might be that the comparison class is (a counterpart of) the individual agent herself. For example, in the case of drunkenness, we sometimes mitigate blameworthiness for agents' obnoxiousness. In such

cases, we're saying that an agent deserves less blame for her obnoxiousness than she would if she were equally obnoxious but also sober. What matters, then, is not that we have two (or more) agents to compare, as is the case in Fig. 1. Rather, what matters is that we have some standard of what is deserved that some particular agent is above or below.

As we can see from the Fig. 1, instantiating the property of being more (or less) to blame is to be understood in terms of the goodness of blame or of the weightiness of desert-based reasons for blame. In such cases, although the desert-based reasons can license the exact same response (i.e., what A and B deserve is the same), they are reasons for that particular response of varying weights.

To see what I mean, an analogy might be helpful. Consider that you and I each have reason to help my daughter. She is a person, after all, and plausibly, we have all reasons to help other persons. But despite the fact that we have reason to do the same thing, it seems that the reasons I have to help my daughter are much weightier than the reasons you have to help my daughter. It would be a mistake for me to forgo helping my daughter for relatively trivial goods. However, it's not so obvious that it would be a mistake for you to do so (you don't, e.g., have to give up date night to buy extra boxes of Girl Scout cookies so that her troop meets their quota). What this shows, I think, is not that you have no reasons to help my daughter, but that mine are considerably weightier. Thus, even though we have reasons to do the exact same thing, given differences in our relationships to my daughter, my reasons are much weightier.

The case of degrees of blameworthiness in Fig. 1 is similar. There, we have reasons to blame two agents for their actions. But our desert-based reasons to blame A are weightier than our desert-based reasons to blame B. This is not because of our relationship to A or B (as was the case in the above example). Instead, it is because of A's and B's respective

relationships to their actions. But questions about what ground these differences (questions that Coates and Swenson (2013), Vargas (2013), and Nelkin (forthcoming) each take up), are questions about the conditions under which A instantiates the property of being more blameworthy than B. So I won't speculate about that here. Instead, I'll simply take Fig. 1 to represent the property that an agent instantiates when he or she is more (or less) deserving of blame and (for now) leave it to others to explain why A and B might be in this situation.


Of course, as I said in the previous section, it's not obvious that this property is really the one that we're typically attributing to agents when we loosely describe them as being "more deserving" of blame. *That* property, as we'll see, is quite distinct.

5. Being Deserving of More Blame

I think that ordinarily, when we talk of people being more (or less) blameworthy or of being more (or less) deserving of blame, we're not making a claim about the goodness of blame or weightiness of desert-based reasons to blame. Instead, we're making a claim about the *amount* of blame that the agent putatively deserves. That is, the "more" (or less) in judgments of "more (or less) deserving" is a quantitative measure. So, what we typically mean when we say that an agent is more blameworthy, is something like "the agent is *deserving of more* blame."

This too can be represented on a Kagan-style desert graph.

Fig 2.



Here, A and B are each represented as deserving blame, since they are in the top left quadrant of a Cartesian coordinate. But as A and B are represented in Fig. 2, neither A nor B is more deserving of blame in the sense of “more deserving” discussed in Section 4. It is equally good from the point of view of desert that A and B be blamed for their actions; i.e., there are equally weighty desert-based reasons for blaming A and B.

However, as A and B are represented in Fig. 2, *what* they deserve is quite different. In particular, the amount of blame that they deserve is quite different. A is represented as deserving twice the blame as B. Now, on a first pass, there’s nothing that’s really unsettling about talk of deserving more blame, but I think that these issues are actually a quite bit more complicated than we typically realize. In particular, I worry that blame doesn’t readily admit of the kinds of quantitative comparisons that are suggested by Fig. 2.

6. Some Complications

As I said, I think something like Fig. 2 represents what most of us are talking about when we say of someone that they are more (or less) deserving of blame. (This isn't exactly right, as we'll see in Section 7, but it's a helpful characterization of one aspect of what's ordinarily at stake.) But on reflection, it's not clear that Fig. 2 really makes sense on a number of leading theories of blame.

Earlier I said that I was going to avoid the question of what precisely blame comes to. Settling that issue has proven difficult, so I don't want to rule out any theory of blame that is *ex ante* plausible (—which is pretty much all of them). But now it seems we're forced to do exactly this. Consider a purely judgmental theory of blame, according to which blame is identified with the moral judgment (or other belief-like state) that another agent has shown ill will in his or her action. If this theory is true, then it's difficult to see what it is that one agent could deserve *more* of than another agent. For suppose that both A and B are deserving of blame. It follows then that A and B each deserve to be judged as having displayed ill will in their actions. What, then, would it mean for A to deserve twice as much judgment that she has displayed ill will in her action? More belief? More confidently held belief? It's just not clear that one can even make sense of the proposal that agents can be *deserving of more* blame on the judgmental theory of blame.

One way to resolve this is to disambiguate a purely cognitive notion of judgment with a form of moral judgment that already includes conative elements. This latter form of “judgment”—the kind of judgment that we're resisting when we tell others “not to judge”—is remarkably close to a number of other theories of blame (including that of Sher, Scanlon, and so-called reactive-attitude theorists). And unlike purely judgmental theories of blame,

Fig. 2 doesn't make complete nonsense of these theories. Take a theory of blame according to which blame is identified with angry emotions like resentment and indignation. On this theory, it makes apparently good sense to say that A is deserving of more blame than B because it makes sense to talk of being more angry at A than B. Emotions like resentment, indignation, and guilt can be felt and expressed more or less strongly. On this picture of blame, then it would seem that what A *deserves more of* is resentment.


But though this makes sense of the thought that A can be more deserving of blame than B, there is another problem with Fig. 2. In the case of criminal punishment, it seems possible to make fine-grained and precise discriminations about what is deserved. Prison sentences are measured in time, and time can be neatly divided however you like. It thus makes a great deal of sense to say of A that she deserves a punishment of twice the length of B. It also makes sense to say of A that she deserves exactly 22.43% more punishment than B. However, it's puzzling to think about how these comparisons apply to interpersonal blame.

It's true, we might sometimes say that we're twice as mad in this circumstance than we were in some other circumstance, but this is, I think, largely for effect. We certainly do get more (or less) resentful, indignant, etc. But it's not clear that this corresponds to discrete quantitative comparisons of the sort one might arrive at if you're thinking about lengths of time. Differences in the amount we blame someone are instead qualitative. I get a little bit resentful. And I also get mad as hell. In between these extremes, I also get very angry and quite indignant. But how much more angry am I when I'm mad as hell than when I'm quite indignant? I seriously doubt there's a way to accurately express that difference in a purely quantitative way.

In Fig. 2 we represent A as deserving twice as much blame as B. Suppose B deserves some resentment. Maybe this means that A deserves fiery resentment. Or maybe A only deserves a lot of resentment, whatever that comes to (does “a lot of resentment” mean that the resentment is especially intense, or that it endures for a long time, or that it has a stronger reasons-giving force, or ...?). And what would half of this come to? Or 22.43% of this? Again, it’s simply not clear. The point here, I want to note, is not merely epistemic. That is, it’s not merely that we can’t ascertain quantitative distinctions when it comes to the emotions. Rather, it’s that there aren’t any distinctions of this sort that can be made, since issues of more or less of an emotional response are not wholly reducible to issues of intensity or duration. Of course, this doesn’t mean we can’t compare two emotional episodes—I can be angrier now than I was yesterday, for example. But it does mean that we should doubt that those comparisons can be made in exactly the way suggested by Fig. 2.

The problem of quantitative fine-tuning doesn’t only arise in the case of emotional theories of blame, but in the case of *any* theory of blame that can make sense of A being *deserving of more* blame than B. It therefore seems that in order to represent a state of affairs in which A is deserving of more blame than B for acting in a type-identical way, we need turn to some sort of qualitative measure. I propose this below in Figure 3.

Fig 3.



In Fig. 3, instead of representing what agents deserve with a point (a “peak” in Kagan’s nomenclature), I’ve represented it as a spectrum that has vague edges. My own view is that this matches our experience of blaming agents more (or less) for their actions than we would have otherwise. There are, after all, no sharp delineations between being indignant and being very indignant. Not even the phenomenal qualities of the emotion can specify this precisely, since in one context, you might be *very* indignant but not attend to your anger carefully, and in another context, one in which you are not distracted from your anger, you might simply be indignant and yet fail to be *very* indignant. In the latter case, you might *feel* more angry, but it might not actually be a case in which you are more angry, since how indignant you are is not simply a function of the intensity of the blamer’s experience. It is


also a function of how that emotion structures your dispositions and future interactions with its target (i.e., the individual who has committed some transgression).

This picture seems to me to be right. For I find it doubtful that the blame an agent deserves can be represented as a precise point. The human responses that constitute blame—whatever it comes to—are themselves too varied to allow for such precision. This doesn't mean that there are no facts concerning how much blame agents deserve for their actions. Rather, it means that those facts must not be more determinate than the practices that give rise to the facts. Fig. 3 respects this constraint and captures the qualitative rather than quantitative way in which we ordinarily take one agent to be *deserving of more* blame than another.

7. Being More (or less) Blameworthy

An agent can be more (or less) blameworthy for her action by instantiating two properties. She can be related to her action in a way that makes it better (or worse) from the point of view of desert that we blame her. This is represented by Fig. 1. She can also be related to her action in a way that makes her deserving of more (or less) blame. Since blaming more (or less) is a qualitative rather than a purely quantitative matter, I take this property to be best represented by something like Fig. 3.

From this, it also seems that an agent can be more (or less) blameworthy by instantiating both of these properties simultaneously. This agent would be more blameworthy if she is (i) *more deserving* of blame and she is (ii) *deserving of more* blame. She would be less blameworthy if she is (i) *less deserving* of blame and she is (ii) *deserving of less* blame. As it happens, I suspect that when an agent is more (or less) blameworthy for an action she typically instantiates each of the properties. Figure 4 represents this.



In Fig. 4, it is better from the point of view of desert that A is blamed for her action. That is, there are weightier desert-based reasons for blaming A than there are for blaming B. It's also true that what it is that A deserves is qualitatively more blame. In other words, not only is it better from the point of view of desert that A be blamed than that B be blame, A also deserves more blame than B.

Earlier I said that talk of being more deserving of blame is typically connected to the property I identified with being *deserving of more* blame. (That is, I think this is the property that a lot of us most readily think of when we think about agents being more (or less) blameworthy.) I stand by that, but I also think that there is a deep connection between

instantiating the property of being *deserving of more* blame (represented by Fig. 3) and the property of being *more deserving* of blame (represented by Fig. 1).

To see this, let's briefly consider how we respond to adolescent wrongdoers and how we respond to adult wrongdoers who are guilty of type-identical transgressions. In the case of adolescents, we typically respond in a milder fashion. This is not mere convention, but instead reflects our sense that adolescents are less blameworthy for their actions than adults who commit actions of the same type. The explanation for this is complicated (this is, in part, what Coates and Swenson (2013), Vargas (2013), and Nelkin (forthcoming) are endeavoring to explain). But surely part of why we typically think of a normal adolescent that she is less blameworthy than her adult counterpart is that adolescents have less self-control over their actions. This is not to say that they have none. Nor is it to say that all adolescents have less control than all adults. Instead, it's just to point out that in typical cases, the truth of the judgment that the adolescent is less blameworthy than the adult hangs on whether she does, in fact, have less control.

But suppose that A is an adult with ordinary levels of self-control and B is an adolescent with levels of self-control that are ordinary for someone at his or her stage of intellectual and emotional maturation. On this assumption it makes sense that we would represent A and B in exactly the way that Fig. 4 does. For presumably, having less control over one's actions would mean that one is less firmly planted in the desert relationship than one would be if one had more control over one's actions. This doesn't mean, of course, that one can't stand in the desert relationship at all; it's just that the weightiness of the desert-based reasons for blaming you will be weaker. Having less control would result in a comparison of the sort represented by Figs. 1 and 4, since in each of those cases, it is better

from the point of view of desert (i.e., there are weightier desert-based reasons) to blame A than B.

The assumption of less control also makes sense of the fact that in Fig. 4, A is represented as being *deserving of more* blame than B. This is because when one has greater control over a wrongful action, it makes more sense that you'd deserve more serious blame for that action. At the limit, an individual would have no control over what she does. But in that case, the individual would not deserve blame for anything she does, no matter how harmful it might be. At the other extreme, an agent would deserve maximal blame if she has total control over her action. If God, for example, could act wrongly, then no one could deserve more blame, since no one else has that kind of control over their actions. This intuition is represented in Figs. 2 - 4, though for the reasons I adduced in Section 6, I think it is better represented in Figs. 3 and 4, respectively.

Fig. 4 represents both of the ways that having less control over one's actions can affect what one deserves. I therefore think that when we speak of someone as being more (or less) blameworthy, it's best not to exclusively think of the property represented by Fig. 1 (as is perhaps most natural given the locution "more deserving of ...") or the property represented by Fig. 3 (as perhaps many of us do) but instead, to think of the conjunction of these properties represented in Fig. 4.

Of course, how blameworthy agents are is not only a function of the control they exhibit over themselves and their actions. Manuel Vargas (2013) very plausibly suggests that an agent's moral ecology also matters for how blameworthy she is. It might be that on reflection, we'll see that this doesn't support Fig. 4, but instead, is better represented by either Fig.1 or Fig. 3 exclusively.¹¹ In other words, although I think Fig. 4 does represent

¹¹ I say "might" because I genuinely don't know.

what it means to say that A is more blameworthy than B when the explanation for why A is more blameworthy is that A has more control than B, this isn't some necessary truth about degrees of blame. Various agential enhancements or maladies that affect the degree to which we are blameworthy for our actions might render us more (or less) blameworthy in only those ways picked out by Fig. 1 or Fig. 3. In any particular case, then, we'll have to carefully attend to the agent(s) in question to evaluate whether and in what ways they are more (or less) blameworthy for their actions.

This suggests an important constraint for those who are working to explain the conditions under which agents are more (or less) blameworthy for their actions. When you propose a way in which having more (or less) control over an one's actions can affect the degree to which one is blameworthy, one must specify the property or properties that having more or less control is relevant to. Or if one grounds degrees of blameworthiness in other features of agents or in their circumstances, then so too, one must specify the precise property or properties one instantiates to a greater or lesser degree. Of course, it might be that any differences will have implications for being *more (or less) deserving* and for being *deserving of more (or less) blame*, but we cannot simply assume this as, I think, Coates and Swenson (2013), Vargas (2013), and Nelkin (forthcoming) have done in their work so far.

8. Conclusion

Blameworthiness comes in degrees. That much seems hard to deny. And yet, what exactly it means to say that A is more (or less) blameworthy than B is not readily apparent. In this paper I have argued that because blameworthiness is a matter of *deserving* blame, we would do well to attend to the question of what it means to deserve more (or less) blame.

Following Shelly Kagan, I use desert graphs to map the possibilities. There are at least three

ways agents can be more or less blameworthy. It can be better to blame them from the point of view of desert; they can be deserving of qualitatively more blame; and it can be better from the point of view of desert that they receive qualitatively more blame. This is what it is to be more (or less) blameworthy.

* * *

As a final note, I wish to simply note that the three properties I've discussed in this paper are not the only possibilities. Kagan's book is incredibly rich and suggestive, and anyone with a cursory knowledge of what he does therein will quickly be able to think of some more properties that, in virtue of instantiating, an agent might be said to be more (or less) blameworthy. For any point that precisely expresses what an agent deserves, Kagan argues that there are corresponding slopes that pick out how good it is from the point of view of desert that an agent receive more or less of what it is that she precisely deserves. Without offering an explanation of this (—Kagan's discussion of these issues ranges over 600 pages, so I hope you'll forgive me for just discussing these issues in such a cursory manner), I'll suggest that if the Western slope of any individual's desert graph is more gentle, if the Eastern slope is steeper, or if both of these obtain, then it will be the case that she is more blameworthy. And by contrast, if the Western slope is steeper, or if the Eastern slope is gentler, or if both of these obtain, then she will be less blameworthy. I've no doubt that many other properties can be identified that render agents more (or less) blameworthy. But for now, we've made a good start: our ordinary practice of holding agents to be more or less blameworthy track normatively significant properties that, thanks to Kagan's help, we can perspicuously identify.