

Injustice and Structural Wrong: *Reply to Sankaran & Monaghan, and Collins*

David Estlund, Brown University

October 21, 2025

Forthcoming in *Ethics*, April 2026

6,900 words

Final edits to accepted version, pre-copyediting

I welcome the chance to reply to two critical pieces about my recent article, “What’s Unjust About Structural Injustice?”¹ which presses a challenge for the view that social structures themselves can be wrong irrespective of wrong by any agents. I hope to show that the argument in that piece does not presuppose that only conduct by agents can be wrong, and also that neither piece provides a satisfactory answer to the challenge: what kind of wrongness is meant if it is not wrong committed by any agent? First it will be helpful to set the stage with a brief sketch of my article’s line of argument, emphasizing a few points that will come up.

I grant for purposes of argument that social structures themselves can be bad in themselves much as natural disasters can be bad, and can be the ground of duties of avoidance, prevention, and remedy. The question is whether they can be wrong. If a certain view about structural injustice does not hold that it is wrong, I say nothing against it. But if, as in much of the central philosophical literature, structural injustice is said or meant to be a kind of wrong and not only something bad, then we might expect it to consist in wrongs committed by agents, along with important social structural implications. For example, in one of the paradigmatic examples of structural injustice, structural racism, culprits include the individuals who wrongfully helped produce the relevant social structure, or who should have done more to prevent it, mitigate it, or do away with it. We could certainly call the structural result a wrong of structural injustice—indeed, this seems often to be what is meant in non-academic political discourse—but there the problematic structure is the fault of some agents. That model would therefore fail to reach some cases that motivate the idea of purely structural wrong, namely, conditions that seem to many people to be wrongs of structural injustice in some way even if they don’t essentially involve or depend on any such wrongs by agents. Examples might be some kinds of hierarchy, or perhaps a regional disparity, or other structural condition that emerges from broad historical causes and continues for some time

¹ “What’s Unjust about Structural Injustice?” *Ethics* 134 (2024) (3):333-359

in a way that is not essentially due to any agentive wrong. Even all required attempts at remedy are hardly guaranteed to succeed. It is reasonable to investigate the intuitive idea that those cases are still wrong in some way.

I argue, however, that the idea faces a deep challenge. Partly in order to cover such cases the most prominent understanding in the philosophical literature about structural injustice, and the one I challenge, follows Iris Young and Sally Haslanger in holding that there is such a thing as non-agentive, purely structural wrong. It's a hybrid view of sorts: the structures are said to be not only bad but wrong, and yet (like mere bads) not wrong by virtue of wrongdoing by agents.² I posed a dilemma for this form of view: If the only wrongs involved are wrongs by agents, then, since, except in few if any cases, social structure is not itself an agent, social structure would not be wrong in itself but, at most, bad. But if there is held to be a kind of non-agentive wrong beyond badness it is hard to see how it would be an appropriate target of what I call grievance attitudes—a certain cluster of reactive attitudes—such as resentment, indignation, or righteous anger. With all of that missing, what is meant by characterizing the structures as wrong rather than merely bad? It's not clear what other marker of wrong can be offered, so we have a hard question. I did not say that these points settle the matter, but I do take them to call for some explanation of what wrongness is when any evident connection to blame is severed (such as warranting attitudes of blame unless there is an adequate excuse). I called this the reach/grievance dilemma that faces that idea of purely structural wrongs of injustice.

Both critical pieces suggest that despite my disclaimers I have foreclosed or disregarded the possibility of any kind of wrong that is not wrong by agents. Kirun Sankaran and Jake Monaghan³ think that my asking, “whether structural injustice involves a distinct nonagentive wrong by seeing whether it warrants” certain individualistic—agent-targeted—reactive attitudes⁴ is “inappropriate,” because [such] “reactive attitudes presuppose that agents are the perpetrators of injustice.”⁵ I don't believe that objection is successful and seeing why requires several steps.

First, they and I emphatically agree that social structure is not the kind of thing that could warrant the “grievance” cluster of individualistic, agent-targeted attitudes around resentment, indignation, and blame where those are understood, as they normally are, as directed at the conduct or quality of will of certain agents. I'll call those classic grievance attitudes, leaving open, here as in the article, whether there might be other grievance attitudes that are not agent-targeted that would fit the case

² Young (*Responsibility for Justice* (Oxford: Oxford University Press, 2011) speaks of “a specific kind of moral wrong, structural injustice, which is distinct from wrongs traceable to specific individual actions or policies.” That's the hybrid view. Likewise, Sally Haslanger clearly states the view that structural injustice—which she interchangeably also calls “oppression” and “institutional injustice”—is a wrong, though “not an individual wrong but a social/political wrong.” See Sally Haslanger, “Oppressions: Racial and Other,” in *Resisting Reality: Social Construction and Social Critique* (Oxford: Oxford University Press, 2012), 311–38, 314 (a lightly revised version of her much earlier essay, “Oppressions: Racial and Other,” in *Racism in Mind*, ed. M. Levine and T. Pataki (Ithaca, NY: Cornell University Press, 2004), 97–123).

³ *Ethics*, April 2025

⁴ Sankaran & Monaghan, 547.

⁵ *Ibid.*, (Abstract), 545.

of purely structural wrong (as, as we'll see, Collins pursues this possibility). I argue that since cases of social structure in themselves could not warrant the classic grievance attitudes we would need some other indication that structures can be wrong and not merely bad. Sankaran and Monaghan agree with the first part, but they see its significance differently. They argue that trying to apply the grievance attitudes to social structures is a hopeless way to inquire into their ostensible wrongness since, as they put it, it would be a kind of "category mistake."⁶ Its being a category mistake doesn't, however, show that test to be misguided, but only that the idea of purely structural wrong is bound to fail it. The relevance of that point is that without that connection to grievance attitudes (classic or otherwise) it is not clear what basis there is for believing there is any such kind of wrong. The form of the dispute is the same as the following: Suppose someone said that it is unclear how a military division can march irrespective of individual soldiers marching, since the division has no legs of its own. Suppose, next, that it is objected that this is an inappropriate challenge to the idea of non-agentive divisional marching (as we might call it), because it supposes that the division is an ambulatory animal, and that is a category mistake. A sufficient reply would be, "Indeed it is a category mistake. And so, it remains unclear how a division, not being an ambulatory animal, could march irrespective of any soldiers marching."⁷

So, despite their implication that this counts against my argument in some way, we agree that the idea of a wrong of structural injustice would be a very different kind of thing from wrongdoing by agents. To address my challenge, then, we want to know what very different kind of wrong the social-structural kind is. They devote quite a lot of space to explaining various important and interesting causal features that can inhere in social structure, and often ones that might at first seem agentively wrong even though it turns out they aren't or needn't be.⁸ Just after explaining two kinds of structural social features ("information costs," and "'invisible' causal structure")⁹ they write, "This shows that even if Estlund is right that social structure cannot be an apt target for reactive attitudes,¹⁰ that isn't a problem for (wrongs of) structural injustice."¹¹ It's not clear, though, why they believe that "this shows" that. If the discussion of specific causal structures is to reinforce the point that social structure in itself could never warrant classic agentive resentment or indignation, then that is not relevant because it is not in question, as I have said. If it is meant, instead, to show that we would fail to understand important things about the workings of social structural inequality if we merely attended to the quality of will of individual agents, then again, that is correct and also not at issue. If it is meant

⁶ Ibid. 546.

⁷ This example is a slight modification, to suit our topic, of an example of Ryle's own when he introduced the term, "category mistake." See, Gilbert Ryle, *The Concept of Mind*, (London: Hutchinson University Library, 1949), Chapter 1, section 2.

⁸ For example, their cases of Schelling's model of segregation (Ibid., 553-54), and "Adam Smith's blistering anti-imperialist polemic against the British East India Company," (554).

⁹ Ibid. 552.

¹⁰ As I'll explain, that overstates my position, but that doesn't affect the present point.

¹¹ Ibid., 553, parenthesis in the original. They mean to signal that on their view it might be redundant.

to show that some social-structural conditions that might seem at first to be due to agentive wrongdoing are not, or at least needn't be, I couldn't agree more.

What is not yet explained is how this is not, as I believe it is, a problem for the idea of purely structural wrong. Now, the following would be a formally valid argument, and some text suggests that they might be tacitly relying on it:

1. There is such a thing as non-agentive wrong of structural injustice that can inhere in social structure itself.
2. But social structures could never warrant classic grievance attitudes since those are warrantable only by agentive wrongdoing.
3. Therefore, there is some species of wrong other than one that warrants classic grievance attitudes, and so other than agentive wrongdoing.

I don't know if they intend this argument, but in any case, the first premise would clearly beg the defining question of my paper—namely, whether there is such a thing as purely structural wrong. We agree, and it is the crux of my own argument, that *if* there is such a thing as purely structural wrong it must not be amenable to classic grievance attitudes such as resentment, indignation, and blame. Far from posing a difficulty for my argument it is my very point.

In a similar worry about having my thumb on the scale, Stephanie Collins¹² suggests that I “appear[] to have judged the matter and decided against the hybrid view.”¹³ But in the text Collins is referring to, I spoke of a “possible wholly deontic conception of structural injustice as wrongful—an extended deontic view.” That would be a view that settles for agentive wrongs, including wrongful failures to prevent, remedy, etc., and regarding any remaining cases of troubling social structure, ones that don't involve or depend on any wrongdoing of any of those kinds, as not wrong but still bad. I continued, in the essay's final sentences, “Does it suffice? It is something of a mishmash of disparate ingredients, seemingly cobbled together *ad hoc*... [leaving] the resulting idea of structural injustice rough and disunified. And nothing we have seen proves that a more unified and deeply structural account is impossible even if it hasn't been developed yet. If it is possible, appreciation of the grievance/reach dilemma may help to point the way.”¹⁴ Some might settle for that extended deontic view,¹⁵ but I stopped short of proposing to do so.

Turning to another issue: Sankaran and Monaghan attend to the provenance of

¹² *Ethics*, July 2025.

¹³ Collins, “Rage,” 707. She says that she will question, “whether our theorizing should take individual human culprits as its starting point—as Estlund implicitly does. I will explain throughout why I don't accept that characterization.

¹⁴ Estlund, “What's Unjust,” 359.

¹⁵ It is, by the way, not different from what I called the “deontic view” itself, defined at p. 341, but the extended formulation exhibits the wider range of cases of agentive culpability that could be invoked.

the term “reactive attitudes” in a famous paper by P. F. Strawson.¹⁶ It is worth pointing out that my argument is not distinctively Strawsonian except to borrow his useful term “reactive attitudes.” Strawson didn’t somehow discover this class of, “attitudes and reactions of people directly involved in transactions with each other; of the attitudes and reactions of offended parties and beneficiaries; of such things as gratitude, resentment, forgiveness, love, and hurt feelings.”¹⁷ I’m happy to use his helpful name to refer to the same rough category of attitudes. “Grievance attitudes,” however, is my own coinage to refer to a subset of reactive attitudes normally associated with attitudes of blame. I do not assume, nor does Strawson indicate that he does, that there are no reactive attitudes other than the several he discusses. Indeed, as I indicated, a good way forward would be to consider whether there are attitudes that are sufficiently like those but which can, without error, be targeted at structures themselves and not at any agent. I’ll call those *structural grievance attitudes*. Humans do often form what appear to be cases of reactive attitudes toward things other than agents, including social structure. The question is whether they could be applicable without error. Consider cases such as being grateful toward your country, or—for grievance attitudes in particular—angry about patriarchal social structure, or resentful of your socially constrained opportunities. To say that those attitudes, when they are aimed at non-agents are not the real McCoy—not real gratitude, anger, or resentment—would be as arbitrary and implausible as saying there’s no such thing as “real” fear at something that is not dangerous. It’s real fear but misplaced. The question I pose is not how structural grievance attitudes could be real grievance attitudes but how, and on what basis, they could be warranted—applicable without error—rather than misplaced. Collins pursues this line, and we’ll turn to her account shortly.

If structural grievance attitudes, such as structural resentment, structural indignation, or structural righteous anger do not presuppose troubling conduct or quality of will then how are they resentment, indignation, or righteous anger at all? And if they aren’t, how is the resented (etc.) social structure plausibly to be counted as not only bad but wrong? It is no argument that some proposed criterion such as this one is unfair or question-begging simply to report that one has no idea how to satisfy it. I take this to be a hard non-rhetorical question, which is not the same as taking it either to beg the question, or to refute or foreclose the possibility of warranted structural grievance attitudes. I made it as clear as I could that I was leaving the question open.¹⁸

The question, as I said, might be answerable by explaining how structural grievance attitudes are enough like classic grievance attitudes to ground a connection to wrong and yet are also not necessarily misplaced even though they are not targeted at agents. Happily, Collins tells us, “This is the strategy I have adopted.”¹⁹ She offers a carefully articulated account of what we are calling structural grievance attitudes, focusing, as do I, on resentment. Before looking

¹⁶ Strawson, P. F. (1963). “Freedom and Resentment,” *Proceedings of the British Academy* 48:187-211.

¹⁷ *Ibid.*, 190.

¹⁸ Estlund, “What’s Unjust,” p. 346, starting with, “I am purposely leaving open the following possible avenue for the idea of culprit-free wrong...”

¹⁹ Collins, “Rage”, 716.

closely at her account, one preliminary clarification: Collins's account posits a certain understanding of social structure, essentially Young's, that she thinks is more amenable than my own to structurally targeted attitudes of resentment. For her, "a social structure is a network of roles (nodes) connected by relations ... [such as] relations of power, authority, and vulnerability... instantiated when enough agents occupy the roles."²⁰ Since my own argument applies to much more than only this Youngian "connection account" of social structure, I defined structure very broadly, counting many things such as, "distributive facts...operations of institutions, a particular legal order, certain pervasive gender roles, widespread norms of family, racial demographics, extent of economic inequality, and average life expectancy."²¹ Rather than try to decide what structure really is I say that, "those who deploy the term or idea get to say, explicitly or by implication, what's relevantly structural (or systemic, or institutional) about the case in their view."²² I do not suggest, as Collins supposes,²³ that I will shape or narrow my understanding of social structure to conform to Young's or any author's specific account, but only that my account will accommodate them. Whatever possibly narrower thing they wish to mean by structure is up to them, and those narrower accounts are bound to be covered. Nor does my broad usage somehow depart from some established narrower usage in the structural injustice literature. Sally Haslanger herself writes in one piece, "As I am using the term here, 'social structure' is a general category of social phenomena, including, for example, social institutions, social practices and conventions, social roles, social hierarchies, social locations or geographies, and the like."²⁴ This is much broader than what Young and Collins have in mind, but social structure can be understood in any number of ways depending on one's purposes. Everything they and Haslanger mention is covered by what I count as structure for purposes of my argument.

So, I don't see any conflict between Collins's narrower and my own broader understanding of what can count as social structure. Moreover, I believe Collins and I are on the same page in several other constructive ways. In order to clearly explain why I am not persuaded that her account succeeds it will be helpful first to lay out certain desiderata that I believe to be common ground between us. This also allows me to lay out much of the shape of Collins's account:

1. The overarching virtue of such an account would be to better anchor the idea that structure itself can be wrong and not only bad, by intimately connecting such wrongness to attitudes that are sufficiently akin to the cluster of attitudes that, so she accepts, are intimately related to judgments of agentive wrong.

²⁰ Ibid., 710.

²¹ Estlund, "What's Unjust," 340.

²² Ibid., "What's Unjust," 339-40.

²³ Collins, "Rage," 707-710.

²⁴ Haslanger, *Resisting Reality*, "But Mom, Crop Tops Are Cute! Social Knowledge, Social Structure, and Ideology Critique," 413.

2. The purely structural wrong involved in the desired account of structural resentment will not be a wrong by any agent, not even a group agent, and not by any of the agents who constitute the social structure.
3. Needless to say, except to highlight the polemical situation, because its goal is to illuminate a non-agentive kind of wrong it will be no defect in such an account that the wrong it posits is not agentive wrong. That would be an irrelevant objection to her account. Likewise, while some views understand attitudes such as resentment as presupposing some quality of will of that which is resented, since we are not foreclosing the possibility of a non-agentive kind of resentment we are not imposing that quality of will constraint on structural resentment. We want to explore whether an account that does not depend on it can succeed without it. The objections I will lay out below do not complain about her account on either of those grounds.
4. I join, in effect, in Collins's positing as a desideratum of anything that is to be counted as a case of resentment that it "respond[s] to slights."²⁵ That is, the content of that attitude type is such that any instance or token of it is misapplied if applied to something that does not commit or exhibit any "slights." "Slights" is her choice of words, and we could substitute others if one wishes, but other natural suggestions for (so to speak) what is resented about what is resented often go straight to presupposing agency, which she hopes to show is not essential to resentment. So, I'm happy to allow "slights" for purposes of exploring that suggestion, though it will turn out to be a crucial question what slights are.
5. By saying, as Collins does, that attitudes of resentment "respond" to slights Collins does not plausibly mean that they never occur unless there was a slight that caused them. Rather, and introducing a term to keep this point clear, I understand her to mean that, as the types of attitude that they are, they, as I shall put it, *presuppose* slights—that, for example, what is resented is resented as having slighted you, just as what is feared is feared as dangerous to you or significant others. Fear presupposes, in this sense, danger. Still, a person can genuinely fear something—say, a spider, or standing at a great height—even if it is not in fact, or not even believed to be, dangerous.²⁶ What fear presupposes is a different question from what, in a given case, is feared (as dangerous). With Collins I'll speak of the "target" of any given instance of resentment to mean that which is thereby resented. A token attitude of resentment might target (by resenting it) something or someone that has not (and might not even be believed to have) slighted you, and this would involve a mismatch between the attitude

²⁵ Collins, "Rage", 711.

²⁶ When it is not dangerous that is one kind of mismatch, and when it is not believed to be dangerous that is a different kind of mismatch, and either can be true without the other, or together. For clarity I will continue to mark the two kinds of mismatch, though I don't see that any issue about this is germane for our purposes.

type's presupposition and the target of a given instance of that attitude.

6. I'll follow Collins in reserving the language of "blame" for conduct of one person expressing to others certain grievance attitudes. Sometimes "blame" or "blaming" is meant to refer directly to certain attitudes, but we can call those grievance attitudes rather than attitudes of blame (including especially resentment) in order to keep clear when we are talking about acts and when about attitudes. That is, I'll here follow Collins in using the terminology of blame to refer to acts. The others to whom the attitude is paradigmatically expressed I will follow her in calling the "audience" of such acts of expression. Just to avoid a certain potential point of confusion (but not any confusion on Collins's part): there needn't be any actual admiration in order for someone to, as we say, express their admiration. Similarly, then, one person "expressing" resentment to another person doesn't settle that the speaker does in fact have the attitude of resentment they are purporting to have. An act of blaming purports to convey that one has a certain attitude, such as resentment. Collins doesn't bring up this point, and I see no tension between it and her account.
7. At least some, maybe all, of the cases of social structure that are meant to count as wrong in themselves are not in themselves agents, and so the needed account will not simply expand the idea of agentive wrongs to include group agents. Collins accepts that social structures "are not 'agents,' and do not have 'agency.'"²⁷
8. Something will be needed to show that the described attitude, despite its fundamental difference from classic agent-targeting resentment, is properly counted as a case of resentment at all. Unless this is satisfactorily shown we would not be given adequate reason to understand the attitude's expression as blame, nor, in the absence of either of those, the proper target of the attitude counting as wrong. Collins appears to agree and offers her further argument (discussed further below) that structural grievance attitudes and their expression would have the same "point" as practices of classic agent-targeted blame, and that on this basis they are "similar enough."²⁸
9. I wrote, and Collins accepts this at least for purposes of argument, that, "It is reasonable...to insist that if a given attitude could be warranted by an unpreventable natural disaster, then it is disqualified as a grievance attitude in the way that matters."²⁹

²⁷ Collins, "Rage," 715.

²⁸ Ibid., 716.

²⁹ Ibid., 715. Estlund, "What's Unjust" 347. Quoting that very passage, Collins misdescribes my view as holding that, "the kind of anger experienced by a victim of unpreventable hurricane damage is not a real grievance attitude." (715) But as I say in that passage my view is that nothing that is really a grievance attitude could be warranted by an unpreventable natural disaster. Given Collins's own view about the connection between grievance attitudes and slights, I don't see how she could disagree. Likewise, (as I discuss throughout the present piece) no attitude that could be warranted

With those desiderata as common ground, Collins proposes the following account of structural resentment: While structure can't be resented for its conduct or quality of will, the attitude's standing as a case of resentment can be adequately established by the role the attitude plays in social practices of some people expressing to others certain negative attitudes about the social structure. Drawing on recent literature, she says that the point of acts of blame might be, "to increase alignment of moral understanding, to scaffold each other's moral sensitivities, to facilitate shared knowledge about normative changes after a wrong, to protest against the target, [and/or] to signal the blamer's commitment to the violated norm."³⁰ Meeting one or more of the proposed "points" can, she says, be said to warrant a grievance attitude.³¹ In summary: in light of the role of the posited structure-targeting attitudes in practices of expressing them to others, the point or points served by such a practice is, she argues, similar enough to agentive resentment and blame in order to count the structures at which they're paradigmatically targeted as cases of purely structural wrong.

A key question for Collins's account is whether it gives us good reason to count what would be structural resentment in Collins's account as a case of resentment, since that is her rationale for counting its targets—certain kinds of social structure—as a kind of wrong. As I understand Collins's the account it invokes the assumption that grievance attitudes (agentive or not) "respond to slights," precisely in order to acknowledge and purport to meet this constraint. When Collins says that, "grievance attitudes... respond[] to slights,"³² she treats the term as self-explanatory. Its usual meaning would be cases of insult or offense in conduct or attitude, as any dictionary will confirm. However, if slights are insults or offenses in conduct or attitude then since structural resentment (by definition) resents social structure, which is no agent, structural resentment would be always, necessarily, misapplied.

Needless to say, this is not the attitude Collins is after. What is resented in structural resentment must instead be, not surprisingly, non-agentive aspects of (especially dynamic ongoing, as is her focus) social structure itself. Presumably, these will consist in features that typically figure in accounts of structural injustice, including, for example, patterns of (static or dynamic) inequality, whether in distribution of certain kinds of advantages, or in asymmetry in certain forms of power, status, or prerogative. In that case the question becomes what can be said to show that such non-agentive structural features should count as slights. That is important, recall, so that the negative attitudes targeting them should count as resentment, and expressing those attitudes in certain ways should count as blame.

by something that is not and is not believed to be dangerous is fear. But just as there is certainly such a phenomenon as misplaced fear, there may very well be a phenomenon of misplaced (and in this case necessarily misplaced), but still real, structural resentment as well. That would depend on whether it is an attitude that resents structure as slighting someone. If some attitude doesn't, I agree with Collins that it is not resentment at all.

³⁰ Ibid., 712.

³¹ Ibid., 711-12.

³² Ibid., 711.

For this, it isn't enough to argue that such attitudes and their expression have similar social effects to what is uncontroversially counted as blaming. That's because nothing blocks the possibility that even some things that are not cases of blaming might serve the same purposes as blaming. Collins seems to anticipate a similar objection: "In case this sounds like I am conflating warrant and consequences—as if grievance attitudes are warranted just when they are useful—note that I am not concerned with any-and-all consequences of grievance attitudes. My suggestion is that a grievance attitude is warranted when it serves the specific point of blame, since grievance attitudes are core constituents of blame."³³ But certain specific consequences are still consequences and no further kind of warrant, and that's all my challenge relies on. Her view is evidently that my resenting you for lying to me is (not only real resentment, but also) warranted so long as, and in the sense that, my resentment contributes to one or more of those kinds of social consequence—whether or not you have, in fact, lied to me. That is, she counts my resenting you as having slighted me as warranted even if you haven't slighted me and even if I don't believe you have. In ordinary language something can be "warranted" in all sorts of ways by being justified or called for, and that can sometimes depend on expected consequences. But this is not the sort of warrant that figures in my argument and it does not address the possibility of a mismatch between the presupposition in a certain attitude type, and its target in a given instance, as when something known to be safe is still feared.³⁴ Salient mismatches here would be resenting someone for a slight that they did not commit, or resenting a social structure even though that is not the kind of thing that can slight anyone. It's true that if Collins's proposed attitudes and their expression did serve similar purposes to classic blame (and I can allow it for the sake of argument) that would indeed show something important if, apart from the present issue there were otherwise strong reason to count these as blame and resentment. But we are not in that position if we don't yet understand how the cases of social structure themselves, which it is agreed are "not agents" and do not exercise "agency," ought to count as slights or slightings in the needed way.

As I argued in my paper, and we've seen that Collins appears to agree, "It is reasonable...to insist that if a given attitude could be warranted by an unpreventable natural disaster [such as a hurricane], then it is disqualified as a grievance attitude."³⁵ So, how does a non-agentive thing like social structure warrant any kind of grievance attitude whereas hurricane damage does not? This brings us to a striking feature of her account, which might be present for this very reason, namely that Collins argues that even though social structure is no agent and does not exercise agency, it is nevertheless "'agential'...in a broad sense," unlike unavoidable natural disasters:

³³ Ibid., 711.

³⁴ See my note 7 and surrounding text, in "What's Unjust" at 336. For fuller discussion of related ideas and important literature around this kind of warrant or fittingness, see, Scarantino, Andrea and Ronald de Sousa, "Emotion", *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.) (Summer 2021 Edition).

³⁵ Estlund, "What's Unjust," 347.

[H]ave our grievance attitudes been downgraded to the frustrated anger of a hurricane victim? No: again, it is crucial that (1) the structure is constituted by agents and (2) the features of the structure that give rise to the grievance attitude are reproduced by the conduct of those agents, however constrained and nonculpable that conduct is. These two facts distinguish social structures from unpreventable natural disasters, squirrels, and stones. Social structures are agentially reproduced, by creatures who are capable of engaging in the Strawsonian participant stance. This makes social structures “agential” in a broad sense, even though they are not “agents” and do not have “agency.” This is unlike hurricanes. Thus, the hybrid theorist can agree with Estlund that grievance attitudes must target objects that are “agential.”³⁶

This, in my experience, is a set of observations that are often roughly invoked in efforts to explain how blaming social structures themselves, but not hurricanes, makes sense, so it’s helpful to have it laid out not roughly but in a clear explicit way. If being what she calls “agential” were something more than being made up of agents and agency in certain ways, but still not a case of an agent or agency, we would want to know what being “agential” is. But her argument is not like that, and she has told us all there is to tell about it. For her, “agential” is merely a word to be understood, “in the broad sense of being constituted by agents and having features that result from those agents’ exercises of agency.”³⁷ But the question is not whether social structure is “constituted by” or “reproduced by the conduct of” of things of a kind that can slight people, but whether it is, even without being an agent, itself the kind of thing that can slight people (presumably non-agentively in some way that would need explaining)—which would be something more. In this passage she has introduced no more, and so the challenge remains and can be reformulated accordingly: that, is of course, all true, but how is it meant to show that structures themselves, and not only the agents that constitute and operate in them, can count in any legitimate sense as slighting? I do not understand her to mean only that we can loosely call them slights; of course we can, and we can loosely call a buggy software app obnoxious or clueless. But then all the argument would show is that we could loosely call structures wrong. That would decline to address whether they are wrong.

Consider a different context in which a parallel point might be made: Suppose you and I agreed that adult human life normally comes with its own moral standing such as rights and duties. But, you argue, there is a kind of thing that counts as adult human life even though it is not an individual human, namely a group of living adult humans. So, you say, groups have rights or obligations too on that basis. Suppose I complain, “Well, I’m open to the possibility of a good argument that a group could have its own moral standing in virtue of some or other connection to

³⁶ Collins, “Rage,” 517. She puts the word “agential” in quotation marks, but she does not appear to be quoting anyone. Her next sentence (again with the scare quotes) is, “Thus, the hybrid theorist can agree with Estlund that grievance attitudes must target objects that are ‘agential.’” This is surprising since I not only didn’t use the term “agential,” I also do not make any use of the broader idea she is using that word to refer to.

³⁷ See also, *Ibid.*, 712.

literal human life, but how does being a group of human lives show that?” In reply, you explain that a group of adult humans is constituted by adult living humans and has features and consequences that result from its being so constituted. Why, you ask, isn’t that a way in which it is, if not literally an instance of adult human life, human-life-like in a broad sense? A good answer by me would be that it is indeed human-life-like in that broad and obvious way, but not in any way that has contributed anything to the question how it could, itself, have its own moral rights and duties. The language of what Collins calls broadly “agential” should not distract us from the fact that no argument is thereby offered that social structure is importantly like an agent, or has agency to some reduced degree, or anything of the kind, nor does she claim otherwise. The fact that structure bears a certain important relation to agents—it is made up of them—fails to explain any legitimate sense in which people can be “slighted” by structure itself even if by no agent at all. (It might be tempting to say the structure’s actions are made up of the members’ actions, but that is not available. It is central common ground that the structure does not perform actions at all, only the members do.) Unless the idea of the structure itself slighting people were shown to make sense it’s hard to see how resenting or blaming social structure itself makes sense, whatever social purposes the referenced practices might happen to serve.

Briefly, we should acknowledge that Collins goes on immediately after that passage to suggest that even if this argument of hers were to fail (as I have argued that it does) it is worth considering whether it might not be needed in the case of some other grievance attitudes, and they might be enough to vindicate purely structural wrong. I say “might” because she says only that this possibility should be considered, and that I had not considered it, which is correct. Even if “agential” in her broad sense is not enough (as I have argued it is not), it might be enough, she says, that structures themselves are “liable to grievance attitudes such as disdain, contempt, disappointment, disapproval, irritation, and criticism.” Social structure might be liable to those attitudes, she says, because they might have “characters” in the right way. She writes (innocently engaging in loose talk just to sketch this possible direction) “they do embed cultures, ideologies, practices, and conceptual schemas that look an awful lot like characters.” Maybe so, although in the cited book by Shoemaker in which a person’s character is linked to possible warrant for disdain, contempt, etc. even if certain other reactive attitudes would not be warranted toward that person, there is no contemplation there of a non-agent, or even any non-person, having a “character” in the sense in which he understands it. His, “tripartite theory of responsibility, [is] a pluralistic approach to the nature of quality of will.”³⁸ Perhaps it could be shown by some other account how relevant kinds of disdain or contempt might be warranted toward the non-agent of social structure itself, and in that way tie structure to wrongness, but we can cross that bridge if and when we come to it.

Finally, I anticipate some readers still thinking, my protestations notwithstanding, that I am illegitimately assuming that nothing is a case of slighting unless it is a case of agency. I am certainly not prepared to explain how it could be anything else, but nor am I insisting on any such thing. To show otherwise would

³⁸ David Shoemaker, *Responsibility from the Margins* (Oxford, GB: Oxford University Press, 2015) viii.

require showing that my arguments to invoke, even if tacitly, that premise, and as I've explained, I don't believe that they do. If someone has a suggestion of how non-agentive slighting could work, I am open to considering it just as I have considered Collins's proposal in detail rather than dismissing it as false by definition. It is Collins who attempts to thread that needle, to suggest that social structure can be a slighting kind of thing. Someone could try some gambit that makes no use of the idea of slighting, but Collins, for reasons I entirely approve, has granted that resentment toward social structure, in order to count as resentment at all, must resent social structure as slighting people. And since, as she grants, social structure itself does not have agency this sets the table for an account of structural slighting that is not a case of any agent's conduct or quality of will. But her strategy, still to connect structural slighting to agency in some way—something she calls “agential”—appears stranded, and the conundrum is right at the heart of the matter. The initial puzzle was how anything could be wrong without being an agent. The conundrum for her appeal to the “agential” is this: structural slighting must somehow be enough like agency (this is her own understandably chosen strategy) and yet cannot be an instance of agency, nor does structure's merely being made up of beings that have agency show structure to be like agency at all—any slighting kind of thing. Her introduction of the broadly “agential” is not an optional detour in her account. If, as I have argued, it fails to provide the intended connection to the idea of slighting, it does not provide the intended connection to resentment, or blame, or therefore wrong, and there's nothing else in the account that does so either.

None of this proves that no account of purely structural wrong can be provided by way of a satisfying account of purely structural grievance attitudes. Collins's proposal takes the bull by the horns, and if I'm right that it fails this may help bring into focus how very sharp the grievance horn of the reach/grievance dilemma is.