

Penultimate Draft

To be published in *An Introduction to Implicit Bias: Knowledge, Justice, and the Social Mind*, eds. Erin Beeghly and Alex Madva. Routledge.

Please cite published version.

Chapter 11

Explaining Injustice: Structural Analysis, Bias, and Individuals

Saray Ayala-López and Erin Beeghly

Why does social injustice exist? What role, if any, do implicit biases play in the perpetuation of social inequalities? Individualistic approaches to these questions explain social injustice as the result of individuals' preferences, beliefs, and choices. For example, they explain racial injustice as the result of individuals acting on racial stereotypes and prejudices. In contrast, structural approaches explain social injustice in terms of beyond-the-individual features, including laws, institutions, city layouts, and social norms. Often these two approaches are seen as competitors. Framing them as competitors suggests that only one approach can win and that the loser offers worse explanations of injustice. In this essay, we explore each approach and compare them. Using implicit bias as an example, we argue that the relationship between individualistic and structural approaches is more complicated than it may first seem. Moreover, we contend that each approach has its place in analyses of injustice and raise the possibility that they can work together—synergistically—to produce deeper explanations of social injustice. If so, the approaches may be complementary, rather than competing.

1. Individuals & the social, in broad strokes

To illustrate the individualistic and structural approaches and how they differ, we'll start with two examples.

Lisa quits her job (adapted from Cudd 2006, discussed in Haslanger 2015)

Lisa is a middle-class woman in a heterosexual monogamous relationship with Larry. They live in a community with expensive childcare and a gender wage gap (i.e., men tend to be paid more than women, in some cases even when they are doing the same work in the same jobs). When they have a baby, Lisa quits her full-time job. One way to make sense of this outcome is to say that there is something about Lisa that makes her quit. For example, it could be that Lisa prefers to take care of the baby full time, or that she is determined to exclusively breastfeed and that requires staying at home. Perhaps Lisa even had a “transformative experience” (Paul 2015). Before becoming a parent, she might have valued her job and planned to keep it. However, maybe the experience of holding a baby in her arms and being the main caretaker for that small being has given her new knowledge about herself and what she really wants. The experience has changed her, let's suppose, to such a great extent that she no longer cares that much about her job and prefers to quit.

Another way to explain why Lisa quits her job is to look at the social system of which Lisa is a part, and understand the outcome as the result of the constraints this system imposes on Lisa. For example, in her society, being a woman positions her as someone with a lower salary compared to her male partner. Besides that, there is no affordable childcare, and babies cannot take care of themselves. All this imposes constraints on what Lisa can do: she cannot keep her job, have her partner quit his to care for the baby, and at the same time keep the most important part of the family income.

Pau tries to communicate their gender identity

Pau is trying to communicate their experiences to friends, in particular, their not identifying as either a woman or a man. Pau says things like “I don’t feel comfortable in public restrooms, I wish there was a non-gendered one I could use.” Their friends take Pau to be confused. They say, “Pau, you are making no sense, maybe you are homosexual, maybe that’s it, but you have to be either a woman or a man.” If we overheard this conversation, we might ask, “What’s gone wrong?” One way to make sense of the problem appeals to the friends’ beliefs and values. Perhaps Pau’s friends are prejudiced against agender people, or trans people more generally, so they interpret Pau’s statements as expressing confusion. Perhaps their binary assumptions prevent them from understanding what Pau has to say, namely, that their gender identity is non-binary.

A second way to understand what’s happening appeals to the wider social environment in which the conversation takes place. Suppose the exchange occurred at a dinner party in the 1990s in Barcelona. The right concepts for interpreting Pau’s experience may not have been available at that time. Though the gay liberation movement had been ongoing for decades and everyone knew what it meant to be “gay” or “lesbian” or “bisexual,” the concept of being “non-binary” was not in widespread use. The concept was missing, in part, because there was no socially acknowledged place to exist outside of the gender binary. Even the concept of “transgender” was largely defined to fit within a binary frame until recently (Stryker 2008). If so, the reason for the distorted interpretations of Pau’s friends is not in their minds, but outside: it’s a feature of the social milieu they inhabit.

These contrasting ways of analyzing Lisa and Pau’s situations offer two different pictures of society, and two different approaches to social justice. In the individualistic picture, we have individuals acting and constraining each other’s actions. In the structural picture, there are other elements such as institutions, laws, social norms, shared concepts like being agender, and material features of environments (e.g. the layout of cities, systems of public transportation or health care). Such beyond-the-individual elements are loosely referred to as structural factors.

In the structural picture, we look at individuals through a wider lens. Individuals are understood as situated in networks of relationships within an organized larger whole, i.e., a structure. In particular, structural analysis reveals how particular individuals are positioned in that structure, which we’ll call, following Sally Haslanger, “a node” (2016). Picture a web of social relations where each node is a type of person. (See Figure 1.) Individuals like Lisa and Pau, as members of different social types (e.g., woman, agender) occupy different nodes and, therefore, have different social positions and social roles. When trying to understand something about an individual, the structural approach asks us to look at the node someone occupies, how that node is connected to other nodes, as well as features of the system as a whole.




Figure 1. A social structure, depicted visually. Each node (black dot) corresponds to a social position, and the lines represent relations between nodes.

Using a structural lens, we see how the behavior of any part depends on its interactions with other parts, and is constrained by the state of the whole. People occupy structural nodes corresponding to their social categories (e.g., sex, race, class, gender identity, national origin). Zach, who sleeps on a sidewalk on a piece of cardboard, occupies a node that is constituted by at least the following dimensions: homeless person, man, white, citizen of the country he lives in. That he occupies this node, and that this node is defined along all those different dimensions, is going to affect how he navigates the city, which opportunities for action are and are not available to him, and how others treat him. For example, as a white man and citizen, he has in principle a significant amount of social power. However, as a homeless person, he will lack credibility, be denied opportunities like the ability to use restrooms in cafés, and be perceived in stigmatizing ways, for example, as dangerous. Dimensions of ability, sexual orientation, and gender identity have similar effects on how they position individuals in the social structure and therefore, on how individuals are treated and what they can do.

The structural picture reveals forms of injustice that might escape the individualistic lens. Think about Lisa's decision to quit her job. Suppose we explain that decision as a result of her beliefs and preferences. Nothing there seems to ring the "injustice alarm." The structural picture highlights, however, that there is more going on. It's not just Lisa's beliefs or desires that cause her to quit. Her personal transformation may be no accident (Barnes 2015). Factors surrounding her invite a radical shift in her preferences, making quitting her job after having a baby the most rational decision for her. What strikes the injustice alarm is that for Lisa (and many middle-class,

married women in similar positions), the rational decision is one that keeps them subordinated, for example, by rendering them economically vulnerable and jeopardizing their careers. The structural explanation captures the system of factors affecting the vulnerable social positions women occupy. It also helps us appreciate that independently of their personal beliefs and preferences (which may vary a great deal from one person to another), people situated in similar positions, and therefore with similar opportunities and constraints, tend to act in similar ways.

Structural analysis is revealing in a second, complementary way as well. Think about the example of Lisa and Larry. Cases like this one have played a prominent role in the history of feminism, and they have serious limitations. As bell hooks points out: “While this issue [of being subordinated in the home as housewives] was presented as a crisis for women, it really was only a crisis for a small group of well-educated white women” (2015a: 38; hooks 2015b: 92). Her point is crucial. Working-class women—many of whom are women of color, some of whom may also be undocumented—may not even have paid maternity leave. For these women, work does not provide “freedom” or “economic security”; they are stuck in exploitative low-paying jobs. The choice to stay home with their children would be perceived as the opposite of oppression. It would be a treasured kind of liberty. A structural approach helps us see this. It calls attention to the fact that women who are positioned differently than Lisa—especially in terms of their socio-economic status and race—may be constrained in ways that may or may not overlap with economically privileged women like her. (See Madva, Chapter 12, “Individual and Structural Interventions” for further discussion.)

There are at least two ways to interpret individualistic and structural approaches to social injustice. First, we can treat them as different metaphysical stories about the *constitution* of society. In the individualistic picture, society and social processes are composed of nothing but

individuals and their interactions. This is called ontological individualism. The structural picture goes beyond individuals and adds social structures and elements to the composition of society. Second, we can treat individualistic and structural approaches as offering two different ways to *explain* what's going on in society. Whereas individualistic explanations analyze social processes in terms of interactions among individuals, structural explanations ask us to take seriously the role of groups in the production of social outcomes. They also adopt a more holistic frame, analyzing society as an interconnected system. Taken in the explanatory sense, the structural picture does not have to worry about questions concerning the metaphysical status of social structural factors. The role of structural factors is (merely) explanatory.

Though it may be tempting to portray structural and individualistic approaches as mutually exclusive, doing so distorts the debate. Proponents of each side sometimes characterize the opposing position in an overly simplistic way, turning it into a straw person. Real-world straw persons—namely, scarecrows—are inadequate copies of the real thing. Similarly, when someone's portrayal of their opponents' views or arguments is described as a "straw person," it means that the portrayal is an inadequate copy of the real argument, and does not represent the strongest and most plausible version of the opponents' position. Accordingly, if we said that individualistic approaches explain social injustice by exclusively appealing to individuals' beliefs or preferences and how individuals interact with one another, while an advocate of the structural approach argues that only structural factors matter, it would be easy to defeat either of these extreme positions. In reality, things are richer and more complicated.

2. Implicit bias and social structures: how they might relate.

To see the complexity, consider the nature of implicit bias. At first, it might seem as if the existence of implicit bias gives straightforward support to a strictly individualistic analysis of injustice. Implicit biases are typically thought to reside “inside our heads.” Many are associated with stereotypes. To have a stereotype, psychologists argue, is to have a set of beliefs or associations with a social group. Consider Pau’s friends. They split gender into two and only two categories, “man” and “woman,” and possess a set of associated gender stereotypes. If so, features of their psychology cause them to act unfairly; hence, it would seem, the primary source of societal unfairness associated with bias resides inside people’s heads. Taking implicit bias seriously does not require such individualistic assumptions, however.

Bias as internalized social structure

Biases enjoy a public existence. Cultural stereotypes, for example, exist as controlling images or ideas in wider society (Collins 2000). Consider the image of a young mother breastfeeding her baby, gazing at the child with complete and utter devotion. The image conveys a message: her baby is all she needs, and it completely fulfills her. Stereotypes such as this are found in novels, movies, online articles, in the jokes and stories that people tell, and in the worksheets that children bring home from school. Intuitively, such images are structural for an obvious reason: they are part of the beyond-the-individual factors that need to be analyzed in order to understand the social world. Yet they may be structural in a more specific way as well. Return to the picture of a social structure depicted in Figure 1; the lines connecting the nodes represent social relations. To the extent that stereotypes and other social biases make social relations what they are, they partially constitute these relations. For example, Lisa and Larry’s relationship is mediated by gender norms; their relationship gets its particular nature, in part, from them. These

norms and images are called “controlling” because they—as social structures do – play a role in influencing what individuals can and cannot do, as well as what they think, feel, hope, and expect from each other and themselves.

Social structures have this power to shape individuals’ lives, in part, because individuals internalize them. Think about Pau’s friends. “You are making no sense,” they say, “maybe you are just a homosexual, but you have to be either a man or woman.” Pau’s friends say this because they have absorbed controlling images and ideas that exist in their wider social milieu. Hence we can think of their biases—whether they qualify as implicit or explicit—as a way in which the social structure manifests in them (Zheng 2018). Something similar might be true of Lisa. Perhaps she quits work, in part, because controlling images of motherhood resonate with her. She may be exercising her autonomy when she shapes her life to match the stereotype; however, in doing so, she may also act as an agent of the patriarchy.

What explains why people are so influenced by social biases, including stereotypes and norms? Human cognition, one story goes, evolved so as to facilitate group cohesion and cooperation (Zawidzki 2013; Haslanger forthcoming). If our minds didn’t attune us to our social environment, allowing us to “pick up” group norms and beliefs, our survival as a species would be compromised. Similarly, Lacey Davidson and Dan Kelly argue that the human mind contains innate mechanisms—modules—that allow individuals to perceive and follow a wide-range of social norms, including norms of reasoning, thought, and action (2018). It is no surprise, according to them, that the gender schema adopted by Pau’s friends is pervasive. The human mind is built to facilitate such uptake.

Bias as gerrymandered perception

Implicit bias connects to social structures in a second way as well. In politics, gerrymandering is a way of dividing up voting districts in a partisan way, so as to make the success of certain political parties more likely. According to some theorists, the same kind of thing happens in visual perception (Munton 2018; for more on bias and perception, see Siegel, Chapter 5, “Bias and Perception”).

Look around you, for example: don’t you see many women in submissive positions at work and in their personal lives? Why is that? According to a structural analysis, the social environment with its norms and arrangements constrains the lives of social groups in a systematic way, and this results in many of their members exhibiting certain properties (for example, being submissive). If middle-class and upper-class women in heterosexual relationships like Lisa tend to quit their jobs when a baby arrives, for example, their economic and social power is compromised. If Larry is making all the money and controls access to the family bank account, Lisa might have to politely ask him for permission to spend money. We may see Lisa doing this or hear her petitioning Larry. However, our eyes and ears cannot access the social backstory. All we see or hear is the outcome: Lisa acting submissively and deferentially toward Larry.

This observation points to something troubling. Suppose you implicitly associate women with taking care of children or with character traits like submissiveness. You may have developed these associations, in part, because you look around the world and see that many—if not most—women embody these stereotypes. Similarly, Pau’s friends might see confirmation of the gender binary in their world. “There are just two genders,” they might argue if Pau pushes back, “just open your eyes and look around.” Statistical evidence might be on their side. However, to the extent that evidence is on their side, this is because agender, gender fluid, and

transgender people are not tolerated, and so, too often, are not publicly visible. Moreover Pau's friends are forgetting all the ways in which children are socialized through the binary and, hence, how the gender binary is actively promoted and collectively reproduced. What they don't consider is whether social reality has been gerrymandered—rigged—to make it appear as if social outcomes reflect unvarnished, unconstrained individual choice.

One might object to calling accurate views of groups “biases.” But even true beliefs about groups may “incline,” hence bias, us towards judging individuals by group membership rather than by facts about them as individuals (Antony 2016; Beeghly, Chapter 4, “Bias and Knowledge: Two Metaphors”; cf. Basu, Chapter 10, “The Specter of Normative Conflict: Does Fairness Require Inaccuracy?”). Likewise, habitual ways of seeing and thinking may become “sedimented” in us, making it harder to be open to evidence when we enter new environments in which our views of groups may not be accurate (Ngo 2017; Munton 2018: 22-26: cf. Leboeuf, Chapter 2, “The Embodied Biased Mind” and Greene, Chapter 7, “Stereotype Threat, Identity, and the Disruption of Habit”). Finally, gerrymandered perception and cognition may constitute biases in that they cause us to think and act in ways that promote an unjust status quo.

Bias as a contextual feature of social environments

Now consider a third approach to implicit bias: “the bias of crowds” model. While traditional theories of implicit bias focus on what's going on “inside the head” of particular biased individuals (cf. Johnson, Chapter 1, “The Psychology of Bias: From Data to Theory”), this new model grounds bias “in the culture, community, and immediate social contexts people inhabit” (Payne and Vuletich 2017: 4).

New data motivates the model. By now, millions upon millions of Implicit Association Tests have been completed on the Project Implicit website. This rich trove of Big Data has enabled researchers to study the geographic variability in individuals' biases (briefly mentioned by Brownstein, Chapter 3, "Skepticism About Bias"). Combining isolated individuals' IAT scores and explicit attitudes to study *overall average* social attitudes across regions, researchers are now uncovering more and more correlations between these average implicit bias scores and a range of regional outcomes and patterns. For example, in countries with larger achievement gaps between boys and girls in science and math, people tend to exhibit stronger implicit gender stereotypes associating men with science (Nosek et al. 2009).

In conjunction with these new data, researchers have also found that individuals' scores on implicit bias tests can be manipulated in various ways. For example, an individual's implicit racial biases can shift dramatically depending on whether they take the IAT in a well-lit versus dark room (Schaller et al. 2003). What *context effects* such as this suggest is that the specific IAT score you get says somewhat less about your biases as an individual—and less about what you're really like deep down and over the long term—and more about the thoughts and images that happen to be floating through your head at a given time.

Imagine this scenario. You are a student attending a predominantly white university in the American South. As you walk to lecture everyday, you see Confederate monuments. You perceive the faculty to be largely white. You know that many of your fellow students are financially stressed out, some are even homeless, while others are living in luxury. How might this state of affairs impact your biases? One set of researchers has examined the question. Here is what they found: average implicit bias scores among college campuses are predicted by broader environmental and structural factors such as the percentage of nonwhite faculty on campus, the

presence (or absence) of highly visible Confederate monuments on campus, and the student body's economic mobility (i.e., the percentage of students who grow up in low-income families but eventually become high earners) (Vuletic and Payne 2019). Significantly, implicit biases seem to be tracking salient inequalities and environmental markers of injustice. If so, not all places inspire bias equally. Modern-day regional IAT scores, for example, correlate with patterns of slavery in the US at the dawn of the Civil War (Payne et al. 2019). In counties and states that had higher proportions of slaves in 1860, white residents have stronger *pro*-white implicit biases to this day, whereas black residents in those same areas have stronger *anti*-white attitudes. There is, in fact, a sizeable and growing empirical literature tracing the psychological and material legacies of slavery across American time and space.

If individuals' biases vary with where they live, go to school, or work, we should perhaps think of biases as existing in environments and situations, rather than as existing in individuals' minds. Advocates of this new model thus defend "a context-based perspective... an interpretation of implicit bias as the cognitive residue of past and present structural inequalities" (Payne et al. 2019: 1; see also Murphy et al. 2018).

These three models underscore a crucial point. Though biases exist in individuals' minds, they cannot be adequately understood as cut off from everything else. Each of these three models connects individuals' implicit or explicit biases to their wider social environment. Psychology and structure are intertwined in deep and important ways. To miss this, or ignore it, is to misunderstand the nature of bias.

3. Comparing individualistic and structural approaches: three criteria

Given the interconnection between structural and individualistic elements of bias, it is too simple to say that an approach has to be either structural or individualistic. The language of priority is more appropriate. Individualistic approaches *prioritize* or *emphasize* the individual, and in particular what is inside their mind, whereas structural approaches *prioritize* elements of the social reality beyond the individual (Madva 2016).

How might we evaluate the strengths and benefits of each picture, individualistic and structural? This section articulates three dimensions along which the two approaches could be compared and evaluated.

One comparison is how accurately each approach identifies what is morally relevant in unjust social situations. Call this *the moral relevance criterion*. For instance, Haslanger claims that an individualistic picture focused on implicit bias “fails to call attention to what is morally at stake” when individuals make choices in unjust social conditions (2015: 1). Recall how the structural picture reveals what is wrong in Lisa’s case: not that she cannot make a choice, but the way her choice architecture is constrained.

A second way to compare the pictures is *the explanatory adequacy criterion*. Each approach—individualistic and structural—explains injustice and social inequality differently. Often these explanations are thought to be competing. If so, the question would be this: which one is superior? But individualistic and structural explanations might be compatible. Perhaps we can keep both in our toolbox.

A third comparison looks at the interventions each proposes, and how effective they are. Call this *the practical utility criterion*. When considering interventions, we might have at least three different aims:

(1) reducing, and ideally eliminating, individual negative attitudes and prejudices;

(2) reducing and ideally eliminating, inequalities (e.g. salary gaps, employment and education opportunities); and finally,

(3) reducing social injustice altogether, and ideally attaining a just society.

These three aims are related, but they are also independent in important ways. (For more about the differences between these aims along with examples, see Madva, Chapter 12, “Individual and Structural Interventions.”) As we compare individualistic and structural approaches, it is important to consider what our aims are in order to determine whether an individualistic or a structural approach is more appropriate.

4. Evaluating structural and individualistic approaches

We now have three criteria. In this section, we apply the criteria and see how the two approaches fare.

The moral relevance criterion. Here is the first question. Which approach is better at identifying the morally relevant features of unjust social situations?

According to Sally Haslanger’s view, structural approaches do better. Haslanger criticizes individualistic approaches for dwelling too much on the motives of wrongdoers (2015: 1). If she is right, these approaches ignore the fact that “the asymmetrical burdens and benefits and inegalitarian relationships imposed on groups” constitute “the normative core” of what’s wrong

with racism and sexism (2015: 1-2). These group-level wrongs become visible only through structural analysis.

How well does this objection work? Remember the three models discussed earlier: bias as internalized social structure, bias as gerrymandered perception, and bias as feature of environments. Each model calls attention to deep connections between individual psychology and social structures. Because such approaches intertwine bias and structure, they do not hide how biased judgments and decisions relate to group dynamics and collective harms.

On one hand, Haslanger need not be disturbed by this result. Her view is not that we must stop talking about bias altogether but that “an adequate account of how implicit bias functions must situate [bias] within a wider theory of social structures and structural injustice” (1). To the extent that newer accounts of bias do this, they do not ignore the “normative core” of racism and sexism. Nevertheless, her objection still has merit. Early theories of implicit bias *did* characterize biases solely in terms of individual psychology, and these theories continue to be influential. Such theories ignore collective dynamics and are problematic for the reasons Haslanger notes. Exhibit A is the philosophical literature on implicit bias, which has been disproportionately focused on questions of individual responsibility (for continued reflection on this point, see McHugh and Davidson, Chapter 9, “Epistemic Responsibility and Implicit Bias”).

On the response just given, Haslanger is open to—and even embraces—more complex accounts of implicit bias. But she has another option. Remember Pau’s friends. Imagine a theorist who argues that these friends have internalized widespread gender norms, i.e., parts of the social structure. When this theorist analyzes what’s ethically wrong with how Pau is treated, let’s suppose, they emphasize the ways in which Pau is harmed by their friends’ binary assumptions. At this point, Haslanger might say: “Ah ha, my point precisely! Explanations of

injustice that appeal to implicit bias—no matter how complex—make folks *more likely* to focus on wrongs to individuals rather than group wrongs, even if they don't necessarily do this, and even if the theories, when properly understood, push against that tendency. So the objection holds: explaining injustice via implicit bias prioritizes individual factors and, in so doing, obscures what's most problematic about social biases.”

We have now reached the heart of the issue. The thought is this. If we endorse an analysis of injustice that prioritizes individuals (and especially their mental states), then structuralists like Haslanger think we are encouraging theorists to remain at the periphery of social problems, ethically and politically speaking, rather than getting to their core.

Let's investigate this thought further. Start with the claim that there is *a normative core to social injustice*. For any injustice, there will be a range of harms and wrongs associated with it. Some of these will be group harms and wrongs. If agender people have no place to exist outside the binary, that harms them as a class. Yet individual wrongs and harms are also present. Pau's friends harm Pau, for example, by acting in a way that defends a rigid gender binary. They fail Pau as friends. Pau can be resentful if they are silenced or remain misunderstood because their friends are dismissive. Likewise, if we want to understand what's wrong with Pau being treated in this way, we ought to think about how it affects Pau's wellbeing and in what specific ways. Perhaps Pau becomes depressed and socially alienated. Maybe there is a certain kind of bodily alienation that accompanies their experiences. If so, there is an imperative to pay attention to—and center in our analysis—Pau's experiences as a particular individual.

On this last point, we should note a powerful tradition in social science: critical race theory (Salter and Adams 2013; Delgado and Stevancic 2017). Theorists in this tradition, as well as feminist theorists, excavate and render visible the experiences of marginalized individuals for

insights into how injustice operates. Writers like Frantz Fanon and Iris Marion Young, for example, eloquently explore how bodily and social alienation feels and functions *from the inside* (Fanon 2008; Young 2005). Such theorists foreground their own particular experiences; yet, quite explicitly, they suggest that these experiences are widely shared and reflect oppressive social dynamics (for additional examples, see Lorde 2007). Their methodology pushes back against the idea that one must center social structures—giving them maximum “air time” in one’s analysis—in order to reach the normative core of racism or sexism. These analyses also reveal that individual and group harms are overlapping and inextricable, so much so that it makes little sense to label group harms “core” while relegating all else to the periphery. To understand group harms, we must understand how oppression affects individuals; to understand the wrongs suffered by individuals, we must appreciate group dynamics.

A second worry deserves to be mentioned here. Haslanger’s objection presumes that the normative core of injustice is stable across all contexts. While it’s a common assumption, it ought to be questioned. Suppose that we are trying to build a social movement to support gender equality. Our central concern might be law and policy. Perhaps we push for laws guaranteeing workplace protections for transgender employees. Maybe we agitate for more generous family leave policies or universal, government-subsidized childcare. To justify these policies, we appeal to how they benefit and provide justice to groups *as a whole*. Given our aims, collective benefits and burdens take center stage here—and rightfully so. Yet this might not always be the case. There could be some contexts in which individual wrongs and harm can and should take center stage, if we care about justice.

Think about Lisa. Imagine that, instead of quitting her job, Lisa is fired when she has her baby. Perhaps, in this specific context, individualistic factors such as her employer’s beliefs

about women, as well as his treatment of Lisa specifically, are of central moral relevance. To get justice in court, Lisa's lawyer must prove that her employer fired her *because of* her pregnancy. If Lisa's lawyer ignores what's in the employer's head and exclusively focuses on widespread group dynamics, she will lose the case. Justice for Lisa will not be served. Similarly, if Pau demands an apology from their friends, it would be very odd if the friends apologized only for harming agender people in general. When we tell a friend that we feel wronged by them, we are generally asking for acknowledgement of a wrong done to us, specifically, as an individual friend. Though group harm may be interwoven with this wrong, the fact remains that the injustice was *done to us*.

Examples such as these lead into controversial territory. They warm readers up to the idea that what's most morally relevant in a situation can change, depending on how you are trying to fight injustice. But the examples should worry us, too. Why should a flawed legal system get to dictate what is most morally relevant in cases of injustice, for example? Justice for many people is not served within the existing system, precisely because intentions are given excessive moral relevance. Employers are often smart enough not to leave a paper trail stating their intentions. When Lisa's lawyer litigates as if prejudice were the key factor in wrongful discrimination, it's thus not necessarily a good thing. She plays into a flawed system and may be seen as perpetuating the false view that bad intentions are required for wrongful discrimination. Bad intentions are simply not always the problem. A group of people might be genuinely committed to social coordination and follow their community's norms without any specific mental state that could be said to be discriminatory, and yet, their community's norms and practices might be such that they disadvantage a subset of the group.

These observations suggest that the moral relevance criterion cannot be used to decisively argue for the superiority of structural approaches. It is not obvious that group harm is always the most morally relevant feature of unjust social situations. What is morally relevant in the courtroom may not be of central moral importance when we are engaging in a collective act of protest. Even more crucially, individual and groups harms appear to be interwoven so thoroughly that it makes little sense to lift up group harm as the essential and most important thing in any and all contexts whatsoever. Both kinds of harms matter, ethically. If so, the moral relevance criterion would push towards a more contextual answer to the question: which approach is better? We would need both approaches to understand what's wrong with injustice, and they would be complements.

The explanatory adequacy criterion.

Perhaps *the explanatory adequacy criterion* tells a different story. Our two approaches—individualistic and structural—correspond to two kinds of explanations used by social scientists to explain the social world and make predictions about it.

When two scientific theories offer an explanation of a phenomenon, how do we know which one is better? What are the most important explanatory virtues and vices? Such questions have long been explored by scientists and philosophers of science.

One view is that structural explanations win: perhaps they offer the deepest and most complete explanation of injustice. Consider our three models of implicit bias. All of them locate the sources of biases outside minds and in social environments. When people internalize group norms and stereotypes, their minds take in controlling images from society at large. The same goes if biases exist environmentally, as residues of historical and ongoing inequalities. In each

case, explanatory priority seems to lie in structures, not minds. Individuals only have the biases they do because they exist in particular social milieus.

On the other hand, serious explanatory gaps may remain if we rely on structural explanations alone. Different individuals respond to social norms and stereotypes differently. Some, like Lisa, embrace them. Others, like Pau, resist them. If we use a structural approach alone, we face serious challenges explaining why some individuals embrace conformity, while others do not. Likewise, there is a strong argument to be made that individuals act on structures. The Stonewall rioters, for example, started a movement that eventually changed American attitudes towards homosexuality, queerness, and gender nonconformity (Stryker 2008). They also challenged unjust laws that permitted the brutalization of queer people. Not only do social structures try to shape individuals, individuals shape the structural aspects of social reality. Within a structural frame, we may therefore want to keep an eye on individual actors for various reasons (Beeghly 2019). Sometimes individuals are complicit and act as agents of structure. Sometimes they subvert structures.

The possibility thus arises that we need both approaches. Maybe they are even compatible and can be used together—synergistically—to explain an event that would be less well explained if only one approach were used. Return to the example of Lisa, who quits her job when she has a baby. Maybe the individualistic picture tells us the proximal or *immediate* cause of her decision: Lisa quit her job because she prefers to take care of her baby. The structural picture might tell about *more prior or distant causes*, for example, how Lisa was socialized to think about gender, or how the possible options for her are constrained to make quitting seem rational. Likewise, the individualistic picture would reveal the immediate cause of why Pau's friends failed to accurately hear what Pau says: prejudice clouds their minds. The structural

picture could tell us that what caused those prejudices in the first place (see Kukla 2014), and it would alert us to the fact that there may not be appropriate concepts like “non-binary” available in the context. In this way, the structural approach might be thought of as expanding the scope of individualistic explanations. Expansion here consists in including more causes, by pushing the causal chain back in time, seeking past causal influences that add to the list of factors that result in the outcome. If so, the two approaches would again be complements rather than competitors.

One objection to understanding explanatory expansion in this way is that it mischaracterizes the right role for structural factors. They are not just back-in-time causes leading to specific mental states in individuals, which ultimately produce an outcome. Structural factors play a role not only as distant causes, but are also present at the very moment the outcome we want to explain is happening (Ayala & Vasilyeva 2015; Ayala 2018).

According to a second sense of expansion, we could see the structural approach as zooming out and including (any) external, outside-the-individual factors. However, as Vasilyeva (in prep, 2016) points out, “not any ‘expansion’ of focus counts,” as not all external factors are structural. Suppose, for instance, that Lisa were forced to quit her job because her malevolent sister locked her up at home for a week. If dynamics in Lisa’s family place her systematically in a submissive position relative to her sister, then a structural explanation could be offered for the outcome. However, if there is no such a dynamic and it’s an aberration from how they usually interact, then there is no structural explanation available. Being locked up, and having to quit her job, is just bad luck. Though an external force is to blame, that force is not structural. As a result, we cannot merely think of structural explanations as zooming out to consider all beyond-the-individual factors.

A third sense of expansion is better for the structural picture: the structural lens neither just broadens the scope back in time, nor merely includes any external factors. Rather, it situates the outcome in a network of relationships within a larger whole, identifying how the relationship between the parts and the whole modify the probabilities of certain kinds of behavior within the system. After all, one way to think of social environments and networks of relationships as influencing what we do in the moment is by *changing the odds*. Part of what makes structures hard to see and understand is that they often don't force or require us to act in certain ways. Often structures just make it more or less likely for us to act in certain ways. Structures put their fingers on the scale: making some actions easier and some actions harder, some options more beneficial and some options less. Structures have these effects, moreover, not just on particular individuals, but for lots of other people in similar situations. They change the odds not just for Lisa's choice but for all the people who occupy "nodes" like Lisa in similar networks of relationships. Significantly, this explanation is not necessarily causal. No causal mechanism for Lisa's choice is cited. Instead what we seem to have on our hands is a mathematical or probabilistic explanation. If so, individualistic and structural approaches provide explanations of different kinds: one explanation is causal, the other is non-causal.

Some readers might be worried by this suggestion. How could a good scientific explanation not be concerned about causes? Answering this question is beyond the scope of this essay. Still, one thing is clear. Though causal explanations dominate the sciences, other kinds of explanations flourish in biology, physics, and yet other disciplines (for a wide variety of examples, see Lange 2016). Some explanations are functional; that is, they appeal to the larger purpose of an event or its role in a system as a whole to explain why it happened. Other explanations are mathematical or probabilistic. (See Fagan ms. for a wider list of scientific

explanations and further discussion.) Though some scientists and philosophers defend explanatory monism (the view that there is only one respectable kind of explanation), it's not the only—or, arguably, the best—view. Explanatory pluralists argue that many kinds of explanation are useful and even necessary for science. A related view—called explanatory particularism (Fagan ms)—says that scientific explanations are the specialized products of particular scientific communities. For example, many psychologists produce individualistic explanations that frame explanations of social injustice in primarily individualistic terms, whereas sociologists tend to produce structural explanations. Good scientific explanations, according to explanatory particularism, are ones that enrich our understanding of the world when combined with others. If so, structural and individualistic approaches could potentially work together to promote a more comprehensive, deeper understanding of social injustice.

The practical utility criterion

We now turn to the final criterion for evaluating the two approaches: *the practical utility criterion*. This criterion looks at what kinds of interventions each approach proposes and compares their effectiveness.

Perhaps ironically, the three models of implicit bias we've examined point to structures—not individuals' minds—as the locus for effective interventions. According to the bias-of-crowds model, individuals' biases as measured by psychologists track the presence of past and present social inequalities in social environments. If so, it's pointless to intervene directly on individual minds. To eliminate bias, we must modify social environments (see also Dasgupta 2013). On a second view, implicit biases are controlling images from wider culture, including stereotypes and social norms, which individuals internalize. Since the root of the problem is structural, so is the

solution. If we could get rid of cultural stereotypes and social norms, we would stop them from being internalized. Finally, if morally problematic biases (e.g. associating women and being submissive) accurately track statistical realities in our world that result from unequal social arrangements, we find ourselves with “a pattern of social inequalities that we can and ought to change” (Antony 2016: 185).

These reflections suggest that, ultimately, the structural picture is the most adequate when considering interventions to fix social injustice. However, before readers with structural affinities get too excited, a little cold water must be thrown on the proposal.

There is paradox surrounding structural change. While it’s one hundred percent true that structural approaches offer the most direct route to social change, a stubborn fact remains. Structures do not magically transform. Individuals must change them. For example, if an unjust law is to be abolished, a huge collective effort will have to be made. People will have to call their legislators and voice their concerns. Investigative reporters will have to publicize the ways in which the unjust law is unfair and harmful. Legislators will have to introduce a new law that invalidates or overrides the old one, and they will have to vote on it. Citizens may have to protest if the vote fails. Though the processes by which social norms and cultural stereotypes transform are less straightforward, the same kind of observation holds. Individuals must act if these aspects of reality are to be changed. For that reason, early advocates of gay rights argued that gays and lesbians would have to come out to their friends and family in order to drum up sympathy for structural change. If your kid’s teacher, your brother, your favorite neighbor, or your daughter comes out as gay, the thought went, it would be harder for pernicious stereotypes to dominate the conversation about gay rights. For example, it would be harder for politicians to argue that employers could simply fire gay people if they wanted to because “gay people were deviants.”

Examples such as these suggest that advocates of structural change must also pay attention to individuals and their mental states if they hope to change structures, even if changing structures directly would be in principle a faster and better way to go.

5. Synergies and convergences

In this chapter, we investigated two approaches to injustice. At first, it seemed as if these approaches—individualistic and structural—were competitors. But that is less clear now. They seem more complementary. We invite readers to look for ways the two approaches can work together.

Here is one exciting example. In psychology, researchers are examining how individuals think about and react to structures. For example, Nadya Vasilyeva and colleagues (forthcoming) propose a psychological intervention aimed at counteracting the way people process social structures by promoting what they call “structural thinking.” Structural thinking acknowledges that people occupy specific social positions, within the landscape of opportunities and obstacles shaped by structural constraints. It recognizes that people (including each one of us) don’t just do things because of underlying biological traits or idiosyncratic preferences (Vasilyeva & Ayala 2019; see also Madva, Chapter 12, “Individual and Structural Interventions” on accentuating the situation). Thus, one thing we have to persuade individuals to do is—*think structurally!* The imperative is to interpret each other not just in terms of beliefs and desires but also in terms of structural opportunities and constraints.

A second lesson of our analysis is this: there is no one-size-fits-all solution for injustice. Sometimes structural approaches may fare better, given our aims. Sometimes individualistic approaches will be more effective. In yet other cases, a mixed strategy focusing simultaneously

on individuals and social structures may work best. Our considered conclusion is therefore this: both approaches are likely necessary to explain what's wrong with injustice, why inequalities occur, and how to transform our world (and ourselves) for the better.

Acknowledgements

A big thank you to Alex Madva for his fierce editorial support, guidance, and encouragement. Erin Beeghly would also like to thank Steven Downes, Melinda Fagan, and Joshua Rivkin for their intellectual generosity and suggestions.

REFERENCES

Anderson, E. (2010) *The Imperative of Integration*, Princeton, NJ: Princeton University Press.

Antony, L. (2016) "Bias: Friend or Foe? Reflections on Saulish Skepticism," in M. Brownstein & J. Saul (eds.) *Implicit Bias and Philosophy: Volume 1*, Oxford: Oxford University Press, 157-190.

Archer, M. (1979) *Social Origins of Educational Systems*, London: Sage.

Ayala, S. & Vasilyeva, N. (2015) "Explaining Speech Injustice: Individualistic vs. Structural Explanation," in D.C. Noelle, R. Dale, A.S. Warlaumont, J. Yoshimi, T. Matlock, C.D. Jennings, and P.P. Maglio (eds.) *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, Austin, TX: Cognitive Science Society: 130-136.

Ayala, S. (2016) “Speech Affordances: A Structural Take on How Much We Can Do With Our Words,” *European Journal of Philosophy*, 24 (4): 879-891. <https://doi.org/10.1111/ejop.12186>

----- (2017) “Comments on Alex Madva’s ‘A Plea for Anti-Anti-Individualism: How Oversimple Psychology Misleads Social Policy,’” *The Brains Blog*, March 6. Available at <http://philosophyofbrains.com/wp-content/uploads/2017/03/Saray-Ayala-Lopez-Comments-on-Madva.pdf>

Ayala-López, S. (2018) “A Structural Explanation of Injustice in Conversations: It’s About Norms,” *Pacific Philosophical Quarterly*, 99: 726-748. <https://doi.org/10.1111/papq.12244>

Barnes, E. (2015) “Social Identities and Transformative Experience,” *Res Philosophica*, 92 (2): 171-187.

Beeghly, E. (Forthcoming) “Embodiment and Oppression: Reflections on Haslanger,” *Australasian Philosophical Review*.

Cudd, A. (2006) *Analyzing Oppression*, Oxford: Oxford University Press.

Collins, P. H. (2000) *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment*, 2nd ed, New York: Routledge.

Dasgupta, N. (2013) "Implicit Attitudes and Beliefs Adapt to Situations: A Decade of Research on the Malleability of Implicit Prejudice, Stereotypes, and the Self-Concept," *Advances in Experimental Social Psychology*, 47: 233–279.

Davidson, L. & Kelly, D. (2018) "Minding the Gap: Bias, Soft Structures, and the Double Life of Social Norms," *Journal of Applied Philosophy*. Accessed via early view:
<https://doi.org/10.1111/japp.12351>

Delgado R. & Stefancic, J. (2017) *Critical Race Theory: An Introduction*, New York, NY: NYU Press.

Dretske, F. (1988) *Explaining Behaviour: Reasons in a World of Causes*, Cambridge, MA: MIT Press.

Fagan, M. (manuscript) *Explanatory Particularism*.

Fanon, F. (1952) *Black Skin: White Masks*, New York, NY: Grove Press.

Gigerenzer, G., Todd, P., & the ABC Research Group. (1999) *Simple Heuristics That Make Us Smart*. New York: Oxford University Press.

Glasgow, J. (2016) "Alienation and Responsibility," in M. Brownstein & J. Saul (eds.) *Implicit Bias and Philosophy: Volume 1*, Oxford: Oxford University Press, 37-61.

Haslanger, S. (2015) "Distinguished lecture: Social structure, Narrative and Explanation," *Canadian Journal of Philosophy*, 45 (1): 1-15.

----- (2016) "What is a (Social) Structural Explanation," *Philosophical Studies*, 173 (1): 113-130.

----- (Forthcoming) "Cognition as a Social Skill," *Australasian Philosophical Review*.

Holroyd, J. and Kelly, D. (2016) "Implicit Bias, Character, and Control," in A. Masala and J. Webber (eds) *From Personality to Virtue*, Oxford: Oxford University Press, 106-133

hooks, b. (2015a) *feminism is for everybody: passionate politics*, 2nd edition, New York, NY: Routledge.

----- (2015b) *ain't i a woman: black women and feminism*, 2nd edition, New York, NY: Routledge.

Johnson, I. R. & Fujita, K. (2012) "Change We Can Believe In: Using Perceptions of Changeability to Promote System-Change Motives Over System-Justification Motives in Information Search," *Psychological Science*, 23 (2): 133-140.

Kaiser, C. R., Major, B., Jurcevic, I., Dover, T. L., Brady, L. M., & Shapiro, J. R. (2013) "Presumed Fair: Ironic Effects of Organizational Diversity Structures," *Journal of Personality and Social Psychology*, 104 (4): 1063-1118.

Kukla, R. (2014) "Performative Force, Norm, and Discursive Injustice," *Hypatia*, 29 (2): 440-457.

Lange, M. (2016) *Because Without Cause: Non-Causal Explanations in Science and Mathematics*, Oxford: Oxford University Press.

Lorde, A. (2007) *Sister Outsider: Essays and Speeches by Audre Lorde*, Berkeley, CA: Crossing Press.

Madva, A. (2016) "A Plea for Anti-Anti-Individualism: How Oversimple Psychology Misleads Social Policy," *Ergo*, 3 (27): 701-728.

Massachusetts Institute of Technology. (2011) "A Report on the Status of Women Faculty in the Schools of Science and Engineering at MIT." Retrieved from http://web.mit.edu/faculty/reports/pdf/women_faculty.pdf

Munton, J. (2017) "Perceptual Skill and Social Structures," *Philosophy and Phenomenological Research*. Accessed via early view: <https://doi.org/10.1111/phpr.12478>

Murphy, M.C., Kroeper, K.M., Ozier, E.M. (2018) "Prejudiced Places: How Contexts Shape Inequality and How Policy Can Change Them," *Policy Insights from the Behavioral and Brain Sciences* 5 (1): 66–74. Accessed at <https://doi.org/10.1177/2372732217748671>

Ngo, H. (2017) *The Habits of Racism: A Phenomenology of Racism and Racialized Embodiment*, New York, NY: Lexington Books.

Nosek, B.A., Smyth, F.L., Sriram, N., Lindner, N.M., Devos, T., Ayala, A., Bar-Anan, Y., Bergh, R., Cai, H., Gonsalkorale, K., Kesebir, S., Maliszewski, N., Neto, F., Olli, E., Park, J., Schnabel, K., Shiomura, K., Tulbure, B.T., Wiers, R.W., Somogyi, M., Akrami, N., Ekehammar, B., Vianello, M., Banaji, M.R., Greenwald, A.G. (2009) “National Differences in Gender–Science Stereotypes Predict National Sex Differences in Science and Math Achievement,” *Proceedings of the National Academy of Sciences of the United States of America*, 106 (26): 10593–10597. <https://doi.org/10.1073/pnas.0809921106>

Paul, L. A. (2015) “What You Can’t Expect When You’re Expecting,” *Res Philosophica* 92 (2): 149–170. Accessed at <http://dx.doi.org/10.11612/resphil.2015.92.2.1>.

Payne, B.K. & Vuletich, H.A. (2017) “Policy Insights From Advances in Implicit Bias Research,” *Policy Insights from the Behavioral and Brain Sciences*, 5 (1): 49-56. <https://doi.org/10.1177/2372732217746190>

Payne, B.K., Vuletich, H.A., & Brown-Iannuzzi, J.L. (2019) “Historical Roots of Implicit Bias in Slavery,” *Proceedings of the National Academy of Sciences of the United States of America*, 116 (24): 11693-11698. <https://doi.org/10.1073/pnas.1818816116>

Payne, B.K., Vuletich, H.A., and Lundberg, K.B. (2017) “The Bias of Crowds: How Implicit Bias Bridges Personal and Systemic Prejudice,” *Psychological Inquiry*, 28 (4): 233–248.

<https://doi.org/10.1080/1047840X.2017.1335568>

Salter P. & Adams G. (2013) “Towards a Critical Race Psychology,” *Social and Personality Psychology Compass*, 7 (11): 781-793. <https://doi.org/10.1111/spc3.12068>

Schaller, M., Park, J.H., & Mueller, A. (2003) “Fear of the Dark: Interactive Effects of Beliefs about Danger and Ambient Darkness on Ethnic Stereotypes,” *Personality and Social Psychology Bulletin*, 29 (5): 637–649. <https://doi.org/10.1177/0146167203029005008>

Stryker, S. (2008) *Transgender History*, New York: Perseus Books.

Thomasson, A.L. (2016) “Structural Explanations and Norms: Comments on Haslanger,” *Philosophical Studies*, 173 (1): 131-139.

Vasilyeva, N. (2016) “Structural Explanations: Structural Factors as Moderators and Constraints on Probabilistic Outcomes,” *Poster presented at the 42nd Annual Meeting of the Society for Philosophy and Psychology*.

----- (Manuscript) “Situating Structural Explanations.”

Vasilyeva, N., Gopnik, A., & Lombrozo, T. (Forthcoming) "The Development of Structural Thinking about Social Categories," *Developmental Psychology*.

Vasilyeva, N. & Ayala Lopez, Saray. 2019. "Structural Thinking and Epistemic Injustice", In B. Sherman & S. Goguen (eds) *Overcoming Epistemic Injustice: Social and Psychological Perspectives*, Washington, DC: Rowman & Littlefield: 63-85.

Vuletich, H.A., Payne, B.K. (2019) "Stability and Change in Implicit Bias," *Psychological Science*, 30 (6): 854-862. <https://doi.org/10.1177/0956797619844270>

Young, I.M. (2005) *On Female Body Experience: 'Throwing Like a Girl' and Other Essays*, Oxford: Oxford University Press.

Zawidzki, T.W. (2013) *Mindshaping: A New Framework for Understanding Human Social Cognition*, Cambridge, MA: The MIT Press.

Zheng, R. (2018) "Bias, Structure, and Injustice: A Reply to Haslanger," *Feminist Philosophy Quarterly*, 4 (1): 1-30. doi:10.5206/fpq/2018.1.4.

SUGGESTIONS FOR FUTURE READING

If you'd like to explore individualistic approaches to injustice further, read:

- Katie Steele, “Choice Models,” in eds. Nancy Cartright and Eleonora Montuschi, *Philosophy of Social Science: A New Introduction*, Oxford: Oxford University Press, 2013, 185-207. Steele introduces readers to decision theory, social choice theory, and game theory and examines how these models seek to explain/predict human behavior.
- Cristina Bicchieri, *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*, New York, NY: Oxford University Press, 2016. Bicchieri offers a predominantly individualistic account of social norms, how they are established, and how to change them. An engaging book full of case studies.
- L.A. Paul, *Transformative Experience*, Oxford; Oxford University Press, 2014. Paul argues that the big decisions in our lives like having a child and choosing a profession cannot be resolved by reason and require a “leap of faith.” She says that one’s preferences and views change radically after big life experiences in ways that cannot be reliably anticipated in advance.
- Mahzarin Banaji and Anthony Greenwald, *Blind Spot: Hidden Biases of Good People*, New York, NY: Delacourte Press, 2013. Banaji and Greenwald analyze bias with a heavily individualistic lens, namely, as a property of individuals’ minds.

If you’d like to explore structural approaches to injustice further, read:

- Richard Rothstein, *The Color of the Law: A Forgotten History of How Our Government Segregated America*, New York, NY: Norton, 2017. Using historical documents and interviews, Rothstein explores how law and policy at local, state, and national levels created residential segregation in the United States throughout the twentieth century.
- Eduardo Bonilla-Silva, *Racism Without Racists: Colorblind Racism and the Persistence of Racial Inequality in America*, 5th edition, New York, NY: Rowman and Littlefield, 2018. Bonilla-Silva analyzes story-telling practices and narratives that generate racial inequalities in the United States.
- Kate Manne, *Down Girl: The Logic of Misogyny* (1 edition), New York, NY: Oxford University Press, 2017. Manne distinguishes sexism and misogyny. She argues that misogyny does not require negative attitudes or beliefs about women, much less hatred towards them; rather, misogyny is a feature of environments in which all genders are policed in ways that maintain gender hierarchy.
- Keith Payne, Heidi Vuletich, and Kristjen Lundberg, “The Bias of Crowds: How Implicit Bias Bridges Personal and Systemic Prejudice,” *Psychological Inquiry*, 28, n. 4 (2017): 233–248. <https://doi.org/10.1080/1047840X.2017.1335568>. This is a “target article” about the bias of crowds model that is followed by commentaries and replies to the commentaries.

- Sally Haslanger, *Resisting Reality: Social Construction and Social Critique*, New York, NY: Oxford University Press, 2012. An important collection of philosophical essays which analyzes social structures, social construction, objectification, and oppression related to race, gender, and other social categories.
- Elizabeth Anderson, *The Imperative of Integration*, Princeton, NJ: Princeton University Press, 2010. A rich analysis of racial segregation in the United States, as well as what's wrong with it and how it sustains injustice and inequality.

If you'd like to read the work of theorists who blend individualistic and structural approaches, read:

- Cherríe Moraga and Gloria Andalzúa (editors), *This Bridge Called My Back: Writings by Radical Women of Color*, Albany, NY: SUNY Press, 2015. In this classic anthology from 1981, you'll find essays, poetry, and philosophy that explore the experiences of women of color in order to better understand how oppression functions.
- Audre Lorde, *Sister Outsider: Essays and Speeches by Audre Lorde*, Berkeley, CA: Crossing Press, 2007. Audre Lorde is one of the most celebrated poets of the twentieth century. In this collection of essays, she talks about her experiences as a black lesbian feminist, as well as her attempts to fight structural sexism and racism. See especially the essay, "The Master's Tools Will Never Dismantle the Master's House."

- Iris Marion Young, *On Female Body Experience: Throwing Like a Girl and Other Essays*, Oxford: Oxford University Press, 2005. Collected essays of philosopher Iris Marion Young. Young argues that understanding and fighting injustice requires both structural analysis and the attention to individuals' embodied experiences.
- Michelle Alexander, *The New Jim Crow: Mass Incarceration In the Age of Colorblindness*, New York, NY: The New Press, 2010. Alexander explores how the United States prison system, as well as the legal system and policing practices have been used to target African Americans and keep them "in their place" since Jim Crow laws were abolished in the 1960s. The core mechanism she describes for both police officers and prosecutors is *unchecked discretion plus implicit bias equals unfair treatment*.
- Kristie Dotson, "A Cautionary Tale: On Limiting Epistemic Oppression," *Frontiers* 33(1), 2012. Dotson explores epistemic oppression and argues that popular ways of analyzing epistemic injustice in philosophy are inadequate. Dotson introduces a type of epistemic injustice (i.e. contributory injustice) that can be seen as a failure of both individuals and structures.