Journal of the American Philosophical Association (2023) 1–20 
DOI:10.1017/apa.2023.20

# Inductive Reasoning Involving Social Kinds

ABSTRACT: *Most social policies cannot be defended without making inductive inferences. For example, consider certain arguments for racial profiling and affirmative action, respectively. They begin with statistics about crime or socioeconomic indicators. Next, there is an inductive step in which the statistic is projected from the past to the future. Finally, there is a normative step in which a policy is proposed as a response in the service of some goal—for example, to reduce crime or to correct socioeconomic imbalances. In comparison to the normative step, the inductive step of a policy defense may seem trivial. We argue that this is not so. Satisfying the demands of the inductive step is difficult, and doing so has important but underappreciated implications for the normative step. In this paper, we provide an account of induction in social contexts and explore its implications for policy. Our account helps to explain which normative principles we ought to accept, and as a result it can explain why it is acceptable to make inferences involving race in some contexts (e.g., in defense of affirmative action) but not in others (e.g., in defense of racial profiling).*

KEYWORDS: induction, social ontology, race, social policy, racial profiling, affirmative action

## Introduction

It is natural to wonder how research on the foundations of inductive reasoning could be important to how we live our lives and participate in society. After all, everyone will continue to reason inductively regardless of what such research reveals. That said, a theory of induction might be useful for evaluating inductive inferences relevant to social policy and for clarifying which policies ought to be implemented and why. That clarity can be useful both for policymakers and for those who vote for them.

Consider the following case, which we take to be representative of a type of tension that many people might experience. Imagine someone who takes themselves to be politically liberal. They oppose racial injustice, support the Movement for Black Lives, and feel anger upon learning of another police killing

We are grateful to Annaleigh Curtis, Shane Gronholz, Sherri Irvin, Chike Jeffers, Mark Lance, Kate Ritchie, and Audrey Yap for their comments on earlier versions of this paper. Saba Fatima, Tiffany Gordon, Jamaal Hyder, Christopher Lewis, and Quayshawn Spencer provided helpful suggestions about framing. We also benefited from feedback from audiences at Dalhousie University, the 2015 California Roundtable on the Philosophy of Race, and St. Mary's College of Maryland. Finally, we are grateful to the editor and referees for their insights.

of a Black person. They oppose policing policies like stop-and-frisk that seem to result in such killings. Moreover, they oppose racial profiling in general because (as a good liberal) they think every person deserves to be judged as an individual and not put into a box in virtue of their social group membership. At the same time, they support affirmative action and believe that it is a good method of engaging in reparations for past injustices. And yet, the liberal feels an uncomfortable tension among these commitments. They worry that affirmative action treats individuals generically in the same way that racial profiling does in so far as it does not judge someone as an individual but reduces them to their social group membership. (Note that we are using this example to help illustrate our project, and we are not here assuming a liberal political orientation.) We will argue that although proponents of racial profiling and affirmative action policies both make inferences involving race, there are epistemic and moral problems that arise for inferences about profiling that do not arise for affirmative action. (We will say much more about how we understand 'race' in section 3. For now, we want to clarify that whatever exactly race turns out to be, it is at least in part a social phenomenon.)

The usual way to defend a policy on the basis of empirical evidence is to begin with a description of a statistical pattern. Call this the descriptive step. Next, one makes an inductive inference, extrapolating the statistical pattern from the observed to the unobserved. Call this the inductive step. Finally, one invokes a normative principle, linking the conclusion of the inductive step to a policy. Call this the normative step. For example, some proponents of racial profiling endorse arguments with the following form (see, e.g., Levin 1992; Schauer 2003; Risse and Zeckhauser 2004; Reiman 2011; and Boonin 2011):

> *Descriptive:* Members of race *R* are observed to commit crime *C* at higher rates than members of other races.
> *Inductive:* In the future, members of *R* will continue to commit crime *C* at higher rates than members of other races.
> *Normative:* Crime *C* (or its effects) is harmful, and therefore law enforcement ought to direct disproportionally more attention toward policing members of race *R*.

Any of these steps can be challenged, but the descriptive and normative steps have received the bulk of attention in the literature on racial profiling. For what it is worth, we think that there are strong objections both to the veracity of many widely accepted crime rate statistics (the descriptive step) and to the normative principles employed in standard arguments for profiling (the normative step). (For examples of such objections, see Thomas [1992], Armour [1994], Colyvan, Regan, and Ferson [2001], Lever [2005], Harcourt [2007], Alexander [2010], Thomsen [2011], Zack [2015, especially ch. 2], and Atenasio [2020].)

In this paper, however, we will focus on the inductive step and its relationship to the others. Here is the basic idea behind our approach. When we reason about race, we make inductive projections on the basis of perceived regularities involving social categories. Until we can explain why the regularities in question hold—which

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

requires an understanding of the relevant social background conditions—we are not justified in taking the inductive step.[1] Such explanation helps us to see that some policies *reinforce* unjust background conditions (which thereby create the data used to justify those policies) whereas others *correct* unjust background conditions (and thereby disrupt flawed inductive loops). Thus, an investigation into inductive inferences themselves can shed light on which normative principles we ought to accept and therefore on which policies we ought to accept. For example, consider this argument for affirmative action:

> *Descriptive:* There are statistical disparities in socioeconomic status across races.
> *Inductive:* These disparities will continue if left unchecked.
> *Normative:* Such patterns have negative social effects and ought to be corrected, and one important way to correct them is to redistribute resources toward members of socioeconomically disadvantaged races.

In this abstract form, the argument is compatible with many extant arguments for affirmative action because many are distinguished by their treatment of the normative step. To mention just a few examples, some arguments focus on correcting past injustices (Thomson 1973) or leveling the playing field (Rachels 1978), whereas others focus on promoting social goods (Nagel 1973; see Boonin 2011 and Fullinwider 2018 for overviews of the literature). We will argue that attending to the inductive steps of the above arguments for affirmative action and racial profiling allows us to identify an important asymmetry and dissolve the uncomfortable tension felt by the liberal in our opening example. More generally, we want to explain part of what it takes to reason responsibly on the basis of race because, as Ritchie (2019) argues, doing so is necessary to fight racial oppression.

Although we focus on racial profiling and affirmative action, our central aim is to provide an account of how the inductive step is relevant to the normative step when it comes to policymaking and to reasoning inductively about social kinds more broadly. In doing so we aim to build a bridge between areas of philosophy that might otherwise seem disconnected, exploring the relationship between the metaphysics and epistemology of induction, on one hand, and the social policies that rely on them (either implicitly or explicitly), on the other. While such connections are already somewhat widely appreciated in the literature on induction, our hope is to contribute something new for those thinking about the relationship between induction and social policies as they pertain to social kinds in general and to race in particular.

[1] See Climenhaga (forthcoming) for a general discussion of the relationship between induction and explanation. We are not the first to discuss this relationship in the context of reasoning about race; see especially Munton (2019b) and also Gardiner (2018). Our approach has a lot in common with Munton's, but it expands on her discussion in two main ways. First, we emphasize the consequences of inductive reasoning for policymaking whereas Munton focuses on evaluating the epistemic status of a subject's beliefs. Second, we describe the epistemic problems for reasoning on the basis of social kinds and discuss the metaphysics of race in greater detail.

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

We are not the first to engage in this sort of bridge-building project. For example, the inductive step has received some attention in recent literature on racial profiling. Atenasio (2020) argues that it is irrational to use race as a basis for projecting criminality given that other factors provide much better baselines. However, most of the relevant literature invokes the phenomenon of moral encroachment according to which our moral obligations to people impact our epistemic duties when reasoning about them (see, for example, Moss 2018: ch. 10; Basu 2019a, 2019b; Bolinger 2020; for a critical discussion see Gardiner 2018). To illustrate, suppose that individuals have a moral right not to be judged to exemplify features statistically associated with their (perceived) race. If that is right, the inductive step fails. Indeed, the inductive step might fail for any argument that involves reasoning on the basis of race, including arguments in support of affirmative action. Our approach is different. Although we recommend paying special attention to the inductive step, we do not argue that the step cannot be justified. Instead, we will explain why we think it is ok to reason inductively on the basis of race (or other social kinds) in some contexts but not others.

Here is how we proceed. In section 1, we discuss the distinction between good and bad inductive inferences. In section 2, we explain why there are some special difficulties in reasoning inductively about the social world. In section 3, we argue that reasoning about race inherits these difficulties because statistical patterns involving race are explained by social phenomena. And in section 4, we apply our theory to standard policies of racial profiling and affirmative action, and we draw some general lessons about the relationship between the inductive and normative steps of any empirically motivated policy defense.

## 1. The Distinction Between Good and Bad Inductive Inferences

The quality of deductive inferences is determined solely by their form. Inductive inferences are notoriously not like this. Consider these two inferences (adapted from Goodman 1983: 73):

1. In the past, all pieces of copper have expanded when heated. Therefore, the next piece of copper will expand when heated.
2. In the past, all men in this room have been third sons. Therefore, the next man to enter the room will be a third son.

The first inference seems much better than the second. Why? For starters, consider the predicates involved. There are genuine respects of intrinsic similarity among all and only pieces of copper; the same cannot be said for third sons. Insofar as the latter category has significance at all, it seems to rely on contingent social features of our world. To use Plato's famous metaphor, only the first category seems to carve the beast of reality at the joints (*Phaedrus* 265d–266a). In other words, the first category seems more *natural* than the second. (We will say more about the concept of naturalness in section 2.) In addition, the regularity in (1) seems to be *lawlike* or at least to be derivable from laws under a wide range of background conditions. (For more background on the metaphysics of laws, see Hildebrand

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

[2020, 2023] and Bhogal [2020].) Indeed, it is somewhat difficult to imagine conditions under which copper would not expand when heated. Aside from situations in which copper is under significant pressure, expansion seems inevitable. In contrast, the regularity in (2) seems accidental. Although it is possible to imagine conditions under which that regularity would continue, it is somewhat difficult. (Perhaps the room is reserved for the Society for Third Sons, though that would be surprising given our background knowledge that being a third son has not been socially significant.) In sum, the regularity described by (1) is projectible—that is, it supports inductive inferences—across a much wider range of conditions than the regularity in (2).

Often we want to know whether a regularity is projectible in a specific context, so it will be helpful to say more about sensitivity to background conditions. If we discovered a fundamental law involving perfectly natural categories, we might be able to make projections *without* regard to background conditions. However, most cases are not like this. Many regularities in physics and most (perhaps all) regularities in the special sciences are sensitive to background conditions—see our discussion of (1) above. In the next section, we will suggest that regularities in social contexts are especially sensitive to background conditions. For example, the regularity in (2) seems unstable; it supports inductive inferences only under highly specific social conditions. For the sake of contrast, however, consider the historical importance of the category of being a first son. Regularities involving that category are projectible in many historical contexts. What explains the differences between these two cases involving third and first sons, respectively? Arguably their categories are on a par, and neither generalization seems to be lawlike. Rather, social background conditions explain the relevant differences between these cases. To sum up, lawlikeness, naturalness, and stability across background conditions are linked in straightforward ways: a highly lawlike regularity involving highly natural properties will be projectible for a wide range of background conditions; regularities that are less lawlike or that involve less natural categories will be projectible for a narrower range of background conditions. However, if we want to make a projection in a specific context, we need to pay special attention to the background conditions–especially when we are dealing with regularities that are not especially lawlike or that do not involve highly natural properties.

It is important to notice that in cases in which a regularity is projectible it is possible to explain why the observed regularity holds. A good explanation appeals to laws or causal structures as well as to background conditions (if applicable). In case (1), the laws do most of the heavy lifting. In contrast, in the exceptional situations in which (2) is a good inference, the background conditions do most of the heavy lifting. In both cases, our ability to make a justified inductive inference goes hand in hand with our ability to explain the original regularity we have observed (and to be justified in believing that relevantly similar background conditions apply to the cases in which we are projecting the regularity). If we want to make good inductive inferences, we need to be sensitive to such issues. In the next section, we will elaborate on our reasons for thinking that special care is required when reasoning inductively about the social world.

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

## 2. Inductive Inferences Concerning the Social World

In section 1, we suggested that the quality of an inductive inference is sensitive to the categories it involves. When projecting from regularities involving non-natural categories, we need to be careful. As it turns out, there are two different natural/non-natural distinctions, depending on whether the categories being distinguished are *properties* or *kinds*. (The overlap in terminology is unfortunate but pervasive; we will not try to correct it here.) Both allow for complications when the categories in question are social.

We understand properties as follows. First, objects share a property if and only if they are similar in some respect. Second, some respects of similarity are more natural, more genuine, than others. For example, charge is natural. There is an objective respect of similarity among any two charged particles. In contrast, the property 'book on the third shelf of my office authored by someone whose last name begins with A, B, J, or Y' is non-natural. This property seems artificial and gerrymandered. Pick any two books with this property; they need not have much in common—besides being books, of course, which suggests that *book* is more natural than the property we are discussing. This suggests a hierarchy of naturalness, with perfectly natural properties—properties that imply exact respects of resemblance—at the very top (see Armstrong [1989]; Lewis [1983, 1986: 59–69]; Rodriguez-Pereyra [2002: 50–52]; and Sider [2011] for more on the distinction between natural and non-natural properties). Third, some properties, such as being popular, depend on social phenomena (norms, practices, conventions, and so on). Call such properties social. Some social properties are more natural than others. For example, the ways in which any two people are popular do not seem to be exactly the same. However, we can still grade social properties by their degree of naturalness. To illustrate, being popular is more natural than either being popular prior to 2022 or unpopular after 2022.

Whereas a property is a respect in which things can be similar, a kind is—or is at least loosely associated with—a metaphysically privileged cluster of properties. Consider the clusters of properties associated with kinds such as *electrons*, *wolves*, and *the popular kids*. Being an electron (wolf, popular kid) requires the instantiation of many different properties, because any two electrons (wolves, popular kids) are alike in many respects. Many of our most interesting generalizations concern kinds rather than properties. We can introduce a natural/social distinction that applies to kinds as well. This distinction tracks two sorts of explanations for why properties are clustered together to form a kind, corresponding to two ways in which a clustering can be metaphysically privileged. Consider copper. The cluster does not depend on any socially contingent phenomena. Other kinds, such as baseball player, depend crucially on social norms and conventions. In neither case are we forced to say that the clustering of properties is entirely accidental. But neither clustering is inevitable; had initial conditions been different, there might have been no copper and no baseball players. Insofar as kinds are concerned, what distinguishes the natural from the social is whether the explanation for the clustering of properties makes essential reference to social phenomena.

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

Here are two quick clarifications. First, strictly speaking, our distinction does not entail that social kinds are not natural kinds (see Khalidi 2015). What matters for our purposes is that there are special difficulties in making inductive projections about social kinds whether or not social kinds qualify as natural. Second, some recent work on social ontology focuses on groups instead of kinds as the basic unit of classification (Ritchie 2013, 2015; Thomasson 2019; Epstein 2019). We will use 'kinds' simply because there is a tradition of appealing to kinds to support a theory of induction (Kornblith 1993), but with minor adjustments 'groups' would serve our purposes just as well.

There are four reasons to be careful when making inductive inferences involving social categories, whether the categories in question are properties or kinds. We will focus on social kinds, but with minor adjustments these reasons apply to social properties as well.

*Reason 1:* Social kinds are especially dynamic; they change over time, sometimes significantly. For example, there were proportionally fewer mustaches and tattoos among cool Seattleites in the early 1990s than among cool Seattleites of the 2020s. Why? Trends change. Natural kinds are less susceptible to such changes.

*Reason 2:* Social kinds often admit of vagueness; they do not usually have clearly defined boundaries. For example, it seems implausible that there is a sharp boundary between being popular and failing to be popular. Because social kinds are clusters of properties, there are two loci for both change and vagueness. The social properties involved in the cluster may change, or the clustering itself may change. Natural kinds seem less vague than social kinds. Simpler, more fundamental natural kinds are often characterized by a small number of essential characteristics defined in precise quantitative terms (see Bird 2018: 23 and Bird and Hawley 2011 for relevant discussion).

*Reason 3:* Social kinds are *interactive*. In general, this means that human behavior is sensitive to our classificatory practices. Here are some examples. People may change their behavior upon learning that they are of a kind or that they are classified by others as members of a kind. This is what Hacking (1995, 1999) calls a looping effect. If you learn that your friends do not think that you are cool, you might adjust your behavior. In addition, people can be sensitive to classification without learning that they are classified in a particular way. If a teacher expects her students to underperform, they might be likely to do so even if she never explicitly conveys that expectation. On the other hand, if she believes that her students are able to meet or exceed the expectations she sets for them, they are more likely to do so (Rosenthal and Jacobson 2003). In both cases, the teacher's expectation turns out to be self-fulfilling. (Although the Rosenthal and Jacobson study is controversial, it nicely illustrates the type of looping effect we should consider when reasoning inductively as well as the fact that not all self-fulfilling prophecies are bad.) As another example, suppose you realize that in videoconferences, because of the lighting and angle of the camera, you appear to be frowning. You compensate by smiling more and being more visually supportive. You then habituate those behaviors and actually become more overtly friendly.

Moreover, human beings can change their own behavior in response to a system of classification that classifies *other* people, which can in turn affect the behavior of those

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

classified. For example, members of the Obama administration did not use the term 'radical Islamic extremism' because, among other reasons, they did not want to vilify their Islamic allies in the Middle East (Stengel 2017). Thus, it is not enough to consider changes in response to classification among those classified; we also need to consider changes in response to classification among those doing the classifying and among those not classified. Human sensitivity to classification provides a unique opportunity for human beings to break from previous patterns of behavior, and thus it constitutes another variable that must be accounted for when making inductive inferences in the context of policymaking. (See Hacking [1995, 1999] for further background on looping effects and interactive kinds; see also Appiah [2005: 65–71]. Mallon [2016: 169–77]) argues that looping effects can support induction under the right circumstances. This does not undermine our claim that looping effects introduce a general problem for projection because Mallon's cases require special knowledge of the relevant background conditions—namely, that we are in the sorts of circumstances in which looping effects support rather than hinder inductive success.)

*Reason 4:* The fourth reason to exercise caution when reasoning about the social world draws on the reasons above and is related to our discussion of lawlikeness, natural categories, and background conditions in section 1: typically (though not invariably), regularities in the social world counterfactually depend on regularities in other domains of science, whereas the latter do not depend on regularities in the social world. Changes to physics would often imply significant changes in chemistry and biology as well as changes in the social world. On the other hand, changes in the social world often would not imply significant changes in biology, chemistry, and physics. The result is a loose hierarchy of regularities, sorted by their modal stability, where more stable regularities hold across wider ranges of background conditions. These claims about stability are plausible regardless of what one might think about the idea that scientific disciplines are unified such that some reduce to others. (For an opinionated introduction concerning the unity of science, see Tahko [2021].)

These features of social kinds suggest that we treat regularities involving social kinds with special care. If we want to make projections based on social kinds for the purposes of making policy, here is what we need to do. First, we must understand why regularities involving those kinds hold. This requires us to be able to explain the regularity in question with reference to the relevant background conditions. Second, we must ensure that the conditions that contribute to the explanation of observed regularities apply to the cases within the bounds of our projection. Third, we must be aware that the policies we adopt or even the reasoning we employ might change human behavior—both among those classified and among those using the classifications. Policies can have various effects on subsequent regularities: they can undermine our inductive inferences or make them self-fulfilling, depending on the circumstances. As a result, if we do not take these steps, we will be epistemically unjustified in projecting.

## 3. Reasoning About Race

In the preceding section we developed a framework for thinking about inductive reasoning about social kinds in general. What should we make of inferences

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

involving race in particular? That depends on the nature of race and on the types of regularities we wish to explain. We will begin by quickly dismissing one possible view concerning the nature of race. It is obvious that racial categories are not properties, perfectly natural or otherwise. Race does not pick out a single respect of resemblance, let alone a perfect respect of resemblance. If races exist at all, they are complex categories such as kinds. Thus, if races exist at all, are they natural kinds or social kinds?

To begin, let us consider what it could mean for races to be natural kinds. Andreasen (2007) argues that races can be distinguished by a cladistics-based method of classification, and Spencer (2012, 2014) argues that races can be distinguished by genetic clustering results among sets of population groups. Since Andreasen and Spencer define races in biological terms, their views can be interpreted as treating races as natural kinds. However, classifying races by a mere appeal to cladistics or genetic clustering will not explain the types of statistical regularities featured in arguments for racial profiling and affirmative action. An explanation of those regularities in terms of natural kinds would require the natural kinds in question to be robust in a certain way. Races would need to be characterized by heritable differences in personal disposition or ability relevant to behavior and specifically to the sort of behavioral traits relevant to offending rates or various socioeconomic indicators. But they are not. Such heritable differences simply have not been found–and not for lack of effort. Moreover, studies by evolutionary biologists and geneticists seem to support the conclusion that different races are genetically equivalent in all interesting behavioral respects. (For some recent overviews of the status of race in contemporary biology see Zack [2002]; Taylor [2004: 48–52]; Glasgow [2009: ch. 5]; Pigliucci [2013]; and McPherson [2015: 679–82].) In addition, we think there is an excellent positive case that explanations of these regularities essentially depend on social factors. (For some defenses of the view that races are socially constructed, see Haslanger [2000, 2008]; Taylor [2004]; and Jeffers [2013b]; for a helpful analysis of the role that social factors play in the production of crime rates, see Alexander [2010], especially chapters 3 and 4.) Thus, we believe that any explanation of racial regularities includes indispensable social elements.

For convenience, then, we will assume that races are social kinds. It is easier to say 'races are social kinds' than to say 'certain kinds of regularities involving race can be explained only by appeal to social factors'. But our arguments do not actually require races to be social kinds. In addition to being compatible with the biological realisms of Andreasen and Spencer, our arguments are compatible with the versions of antirealism about race endorsed by philosophers such as Appiah (1986, 1996), Zack (1993), Glasgow (2009), and Blum (2010). What is essential to our arguments is that the explanation of the relevant regularities is partly social. All of the philosophers just mentioned would agree with this claim. (For a similar reason, Mallon [2006] argues that the metaphysical gulf between antirealists and social constructionists regarding race is small.)

As we argued in section 2, regularities involving social kinds can be projectible. However, we must use special caution when reasoning about them. Accordingly, we need to exercise caution when it comes to assessing the inductive step of any

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

evidence-based policy concerning race. To do so, we must be able to give an explanation of a certain sort. Discussing statistical correlations between races and social conditions, Paul Taylor says this:

> The failure of classical racialism means that *because they're black* is no longer an explanation for anything. It becomes, instead, a gesture at a request for an explanation, or for an answer to a question like this: What is it that links black people to these social conditions? (Taylor 2004: 85)

We are building on Taylor's point, arguing that epistemic responsibility not only requests but demands an answer to Taylor's question. If we want to reason responsibly on the basis of race, we need to be attentive to background conditions, both natural and social, and consider various possible looping effects connected to policy proposals. In the next section, we will connect this conclusion to policies of racial profiling and affirmative action and draw some general lessons for policymaking.

## 4. Putting It All Together

Recall that many defenses of a proposed social policy begin with a descriptive step describing a statistical regularity. Two further steps are required. The inductive step projects that regularity from observed cases to unobserved cases. The normative step recommends a policy in response to that pattern. In section 2, we sketched an account to guide us in satisfying the epistemic demands of the inductive step in social contexts. Our central claim was that when reasoning about social kinds, we need to be able to explain *why* a regularity occurs before we project for the purposes of choosing policy. As we saw, providing such an explanation requires special attentiveness to social background conditions and awareness of potential looping effects. In section 3, we argued that regularities involving race are highly dependent on social context and thus that reasoning about race requires the relevant sort of care. In this section, we will apply our account to standard policies of racial profiling and affirmative action. However, we will first describe two conclusions concerning the relationship between the inductive and normative steps of empirically motivated policy defense.

Our first conclusion is seemingly modest but important nonetheless: Investigating the background conditions required to explain an observed regularity regarding social kinds is likely to reveal facts that bear on whether a policy is just or unjust. As our account of induction rationally requires us to be able to explain regularities prior to projection, we thus have a link between the inductive and normative steps of policy evaluation.

Here is a case that illustrates this point:

> Insurance companies that only used actuarial criteria in the assigning premiums in life insurance rating tables in the United States between 1920 and 1970 that were broken out by race, gave significantly higher rates to African Americans. Actuarially, this could be defended to a

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

> State Insurance Commissioner as warranted based solely on actuarial figures. But the statistics do not give the entire picture. The *reasons why* African Americans had a skewed mortality rate against the European-descent population was because of the fall-out of various social problems from the post-slavery era: job and wage discrimination, inferior medical treatment options, poverty and poor nutrition, inferior apartheid living environments, and lynching. These factors were instances of oppression by the mainstream society against African Americans. Are the victims required to pay *twice*? This is unjust (Boylan 2004); thus though the actuarial figures for life insurance show a higher [mortality (corrected from 'morality' in the original publication)] rate, this should not be a legitimating factor for charging higher rates to African Americans—*because it was not their fault that they were victims*. Society has a negative duty to absorb the cost of their higher respective rates, since it was in that society that the harms occurred. (Boylan 2011: 141)

Boylan's analysis nicely illustrates how asking questions about *why* a pattern holds is relevant to the matter of *how* policy makers ought to respond. The briefest introduction to American history pertaining to slavery, civil rights, immigration, and so on should alert us to the fact that explanations of racialized regularities often contain morally relevant information. We are taking the obvious point that history and background conditions matter and applying it to the specific issue of how the inductive step and normative step of a policy defense are related to one another. As obvious as this might sound, it is often missing from the way policies are evaluated in the real world, and thus we see value in making it explicit.

We have established a connection between the inductive and normative steps. We cannot simply observe a pattern, make a projection, and then appeal to our normative principle of choice to support our preferred policy. There is a gap between the descriptive and inductive steps that must be bridged, and the type of bridge we build can inform our choice of policy. New information about background conditions can help us to acquire new concepts, influence our choice of normative principles, or change the application of normative principles we already accept.

Suppose we have good reason to believe that, other things being equal, the observed statistical pattern on which a proposed policy is based will continue to hold in the future because the background conditions that explained the observed pattern persist. Despite all this work, we are not yet in the position to move on to the normative step and decide on policy. Given the interactive nature of social kinds, when considering policy options, our account also requires us to be cognizant of human responsiveness to policy—that is, to potential Hacking-style looping effects (Hacking 1995, 1999). This leads to our second general conclusion: Policies concerning interactive kinds can influence the very conditions required for the inductive step to be epistemically justified.

One especially problematic way in which this can occur is for policies to create (or sustain) the very conditions that are required to make a regularity projectible. An

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

example of this is what Sally Haslanger calls discursive social construction: 'Something is discursively constructed just in case it is the way it is, to some substantial extent, because of what is attributed (and/or self-attributed) to it' (Haslanger 1995: 99; see Langton [2009] for another example of this phenomenon). Suppose that parents believe that girls require fewer calories than boys, that caloric intake plays a biological role in determining how big children grow to be, and that size plays a role in determining how many calories someone requires (Jelenkovic et al. 2016; Khazan 2014; Jaggar 1987: 34). In response, parents feed girls less than boys. Over time, such beliefs can contribute to the outcome that girls in general require less food than boys. With these general conclusions in mind, we can finally apply our theory to racial profiling and affirmative action.

To begin, let us revisit the generic sort of argument that some put forward in favor of racial profiling, and let us suppose that it targets racial minorities in the United States:

> *Descriptive:* Members of race *R* are observed to commit crime *C* at higher rates than members of other races.
> *Inductive:* In the future, members of *R* will continue to commit crime *C* at higher rates than members of other races.
> *Normative:* Crime *C* (or its effects) is harmful, and therefore law enforcement ought to direct disproportionally more attention toward policing members of race *R*.

Our account requires us to explain why there are crime rate disparities before we project for the purpose of making policy. Given the failure of biological realism about race to explain crime rate disparities, the explanation must be social in character. At a certain level of abstraction, the explanation is obvious: there is a long history of racial oppression, the residual and ongoing effects of which lead to massive socioeconomic disparities (among other things); these play a significant role in explaining crime rates. Now suppose that a law enforcement agency adopts a policy that disproportionately directs resources toward race *R*. Looking for crime *C* in community *R* will impact subsequent crime rate statistics: the increased attention on *R* is likely to uncover more instances of *C* among *R*, while the decreased attention on non-*R* is likely to uncover fewer instances of *C* among non-*R*. In other words, if police target a certain group, subsequent crime rate statistics will tend to confirm the background belief that members of that group are especially prone to commit crime, leading to a cyclical effect (what Bernard Harcourt [2007] calls a ratchet effect).[2] In addition, such a policy will tend to lead to the result that members of *R* disproportionately experience the negative effects

[2] See Harcourt (2007: 28–29) for a concrete example. Lever (2005: 97–98) and Jeffers (2013a: 129) criticize racial profiling on similar grounds. Alexander (2010) provides a detailed account of how policies can lead to cyclical effects. Gendler (2011) includes a summary of recent empirical work on human psychological mechanisms that may contribute to such cyclical effects; see also Egan (2011: 75–77). See Munton (2019a) for an argument that visual experience itself can be susceptible to bias by encoding previously experienced regularities; her discussion of 'gerrymandered priors' in section 4 of her article is especially relevant.

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

of the carceral system, further exacerbating socioeconomic inequalities among races. These are exactly the sorts of morally vicious self-fulfilling prophecies we should avoid when making policy. To be clear, we do not mean to suggest that socioeconomic status is the only causal factor in explaining crime rate disparities, but it is important, and because it is familiar, it serves as a good example.

We have examined some cases in which looping effects lead to social problems. However, social problems are not a necessary consequence of looping effects. In a recent paper, Koskinen (2022) points out that the use of 'language nests'—programs designed to promote the use of Indigenous language in Indigenous communities—can help to revitalize Indigenous languages. Such revitalization replaces old patterns with new ones. The goal of such programs is precisely to undermine existing regularities. Under the right conditions, the new patterns concerning Indigenous language use may even become self-perpetuating, in a virtuous way as part of a thriving culture.

Now let us consider a typical policy of affirmative action. We take a typical policy of race-based affirmative action to require decisions about hiring, admissions, and so on, to take race into account in the following way: members of certain historically disadvantaged groups are to be given preferential treatment in some sense. In some respects, policies of race-based affirmative action and racial profiling are analogous: both involve the allocation of 'resources' to various groups, and both can be motivated by race-based statistics. Yet there are important differences. For our purposes, the crucial one is this: the policies impact members of groups in very different ways. To see why this matters, let us revisit the argument for affirmative action we mentioned earlier:

> *Descriptive:* There are statistical disparities in socioeconomic status across races.
> *Inductive:* These disparities will continue if left unchecked.
> *Normative:* Such patterns have negative social effects and ought to be corrected, and one important way to correct them is to redistribute resources toward members of socioeconomically disadvantaged races.

As before, our theory requires us to exercise caution when making the inductive step. We need to explain socioeconomic differences before we can project for the purposes of making policy. The explanation shares part of the same basic structure as the explanation for crime rate disparities: namely, that there is a history of oppression whose residual effects lead to unjust and unfair socioeconomic disparities. We might add that due to features of our economic and political systems, unfair and unjust distributions of wealth are passed on through generations even if the unjust laws from which socioeconomic inequalities originated have been abolished. Obviously, additional factors could be added as well. What matters for our purposes is that: (i) such explanations are social; (ii) in principle, at a certain level of abstraction, we can provide a good explanation of socioeconomic disparities; and (iii) such explanations reveal facts relevant to the normative evaluation of affirmative action.

At this point, recall the third requirement: we must be mindful of looping effects. Here we find another crucial difference between our two cases: whereas policies of

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

racial profiling have a tendency to harm members of historically oppressed groups and thereby *create* the conditions that justify the descriptive step, policies of affirmative action have a tendency to benefit members of historically oppressed groups and *challenge* those conditions. Thus, we have a significant difference between our two race-based policies. We can be perfectly consistent in claiming that typical policies of racial profiling are unjust and that typical policies of affirmative action are just. Moreover, because policies of affirmative action aim to *correct* the underlying social conditions that lead to the regularity in the first place, they are not self-fulfilling in an epistemically problematic way. Indeed, there is a sense in which a successful policy of this sort will ultimately render itself obsolete as the social conditions required to justify the inductive step are changed by the policy.

Before concluding, we will consider a few objections. First, once we have recognized the different effects of the policies discussed above, the inductive step may seem to lose its importance. Why not just examine the effects of policies (focusing on the normative step) and not worry about the inductive step? For starters, we will be less likely to be aware of relevant normative facts. In addition, if we are not in the position to explain the regularity in question, we will not be in the position to predict the effects of the policy accurately (for example, because looping effects will be more difficult to identify). As mentioned above, ignoring the inductive step can lead to harmful self-fulfilling prophecies such as Harcourt's (2007) ratchet effect.

Second, it might appear that we have moved too quickly or that the explanations above are less obvious than we have suggested. As any social scientist would attest, it is difficult to provide fine-grained causal explanations for features of the social world as well as to verify empirically the types of looping effects we described above. Is it really so obvious that the history of racial oppression is responsible for the asymmetry between the types of policies we have discussed?

We think it is. Coarse-grained explanations are appropriate for our level of abstraction. (See Chetty et al. [2020] for a nice example of an attempt to provide more fine-grained explanations of racial socioeconomic disparities. Note: We take their fine-grained explanations to complement, rather than compete with, the coarse-grained explanations we have provided here.) One reason for this is that the arguments for racial profiling and affirmative action have an important similarity. Crime rates—at least for the sort of crimes that garner attention in proposals for racial profiling policies—are strongly linked to socioeconomic status. Thus, the explanation of crime rates (the descriptive step in the argument for racial profiling) and the explanation of socioeconomic patterns (the descriptive step in the argument for affirmative action) are close relatives. In both cases, socioeconomic status plays an important—though not necessarily complete—role in explaining the relevant regularities. Given the significant differences between the policies, we should expect them to have very different effects. As a reminder, our judgments about racial profiling and affirmative action are relativized to a particular context. Under different circumstances or in a different world, there might be no asymmetry between the two sorts of policies. Though our goal might be to bring about a world where the asymmetry in policies would not apply, we do not yet live in such a world.

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

A third objection is that a dilemma arises concerning the relationship between our account and attitudes of risk avoidance. On the one hand, suppose that our account encourages an attitude of risk avoidance on the part of those involved in making policy. Kinney and Bright (2023) argue that risk aversion can make it instrumentally rational for persons of privilege to ignore information that might make them aware of their own privilege. Granting Kinney and Bright's point, if our account encourages risk avoidance, then it would seem to support some of the very attitudes that might sustain the social conditions that we are trying to change! Encouraging risk avoidance also threatens to make our account incompatible with recent work that seems friendly to our project. For example, Bovens (2016) argues that if one wants to hire the best candidate, one should pursue proactive affirmative action policies that allow one to shortlist minority candidates who are less qualified on paper. Doing so is a little risky, but it increases the possibility of hiring the best, and thus the expected payoff is higher. If we encourage risk avoidance, we undercut such arguments for affirmative action. Fortunately, we do not think our account encourages a general attitude of risk avoidance. Our above discussion of the grain of explanation is relevant here. We will never have complete knowledge of the social world; it is too complex. But we can understand it well enough to see that some explanations of regularities are better than others, and when the stakes are high, we should try to explain observed regularities as best we can. In sum, we do not think our project is in tension with recent work on risk avoidance.

We will now consider the other horn of the dilemma. Suppose that our account does *not* encourage an attitude of risk avoidance. This raises a question: Why would it not be reasonable for someone who lacked an explanation of a regularity but was confident that it was not accidental to say: 'I don't know why the regularity has occurred, but I'm willing to bet it'll continue. In fact, I'm so confident that I'm willing to base policy on it!'? We grant that someone could be instrumentally rational in adopting such an attitude, but on our account such a person would be unable to defend the inductive step of an argument for their preferred policy. That person would be guilty of an epistemic error; they could not promote their preferred policy because they would lack epistemic justification for the inductive step of their policy defense. Specifically, they would exhibit a failure of sensitivity to the difficulty of reasoning about the social world: they would fail to apply lesson one, concerning explanation, and also fail to account for lesson two, concerning looping effects.

So much for objections. Let us return to the general lesson of this section: The normative impact of a policy cannot be properly assessed without recognizing its interaction with the inductive step. A particularly nefarious policy may contribute to oppression while misleading well-intentioned people into thinking that it is instead responding to, and perhaps even correcting, a social problem. Though we have not conducted a careful analysis of the specific background conditions relevant to inductive inferences involved in typical policies of racial profiling, it is not implausible that these policies of racial profiling could be nefarious in this way. Thus, we have a prima facie case against them: if we are not careful, we will not be epistemically justified in making the inductive step, and thus these policies

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

of racial profiling will be (epistemically) irrational. If we are careful to meet the demands of our account of induction, we are likely to reveal moral facts that imply that these profiling policies are unjust. In contrast, a virtuous policy does not have the potential to perpetuate injustice. Though we have not conducted a careful analysis of the specific background conditions relevant to inductive inferences involved in affirmative action, it is not implausible that such policies could be virtuous in the way we described, at least insofar as they seem likely to correct, rather than reinforce, the unjust social circumstances that produce the relevant statistics. Thus, we have a prima facie case in their favor. At the very least, we have identified an important asymmetry between policies of racial profiling and affirmative action—an asymmetry that could go unnoticed if one is not attentive to sociohistorical context.

According to our account, then, reasoning about race is not always unjustified, immoral, or unjust; we ought not adopt a policy of epistemic colorblindness according to which we should not ever reason on the basis of race. Although calls for colorblindness are not as popular as they once were, they retain some powerful advocates. Consider what United States Supreme Court Chief Justice John Roberts said in *Parents Involved in Community Schools v. Seattle School District No. 1*: 'The way to stop discrimination on the basis of race is to stop discriminating on the basis of race'. (See also the majority opinion in *Students for Fair Admissions, Inc. v. President and Fellows of Harvard College*.) Although such an approach might be appropriate in some utopian future, in a world shot through with injustice such an approach serves only to perpetuate the status quo. (For two helpful and recent academic discussions of colorblindness see Bright [2017] and Harris [2017].) Different regularities require different background conditions for their explanation. Different background conditions will have different morally salient features. And different policies will have different kinds of (morally salient) effects. As a result, the relationship between the inductive and normative steps will not be the same in all situations. Our account requires sensitivity to social background conditions in which a policy is to be implemented, lest we fail to justify the inductive step. As a result, our account yields different conclusions for different kinds of policies and different conclusions for similar policies implemented in varying social contexts. These are virtues of our account. Although our account begins with an abstract, highly idealized picture of the natural world, it is applied in messy, nonidealized scenarios. It is our hope that in applied ethics, social philosophy, and political philosophy, ideal theorists and nonideal theorists alike will find that our account provides a useful framework for thinking about connections between reasoning, moral obligation, and justice. For this reason, we think of this project as complementing, rather than reforming, much of the extant work on moral issues pertaining to the philosophy of race (see the many citations in section 3), evidence-based policy (e.g., Cartwright and Hardie 2012) and inductive risk (e.g., Douglas 2000).

However, this is not to say that the role of our account is merely complementary. Recall that we began by distinguishing the inductive step from the normative step. On our account, the former has important and underappreciated implications for the latter. Without a refined account of epistemic justification—and, in particular,

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

a theory of how to make inductive projections—and without focusing on the role that inductive reasoning plays in the promotion of injustice, philosophy is likely to tell incomplete stories about injustice. Our account focuses on one of the inner origins of injustice: namely, the methods of reasoning on which proponents of social policies rely. By illuminating the inner epistemic origins of various sorts of injustice, we can provide a more complete account of past, present, and future injustices. This puts us in a better position to effect social progress and promote justice.

A final remark: We have invoked some metaphysics—namely, realism about laws, properties, and (perhaps) kinds, as well as anti-essentialism about race. Is this metaphysical foundation required for our general approach? We are inclined to think so, but we will not defend this claim here. Rather, our focus has been to show how a careful philosophical account of the nature of induction can bear fruit in matters of policy. Those who dislike our metaphysics are encouraged to develop their own theories and show how they can be fruitfully applied to the problems we address. Indeed, one of our hopes is that this article will inspire critics to do just that.

BARRETT EMERICK
ST. MARY'S COLLEGE OF MARYLAND
*bmemerick@smcm.edu*

TYLER HILDEBRAND
DALHOUSIE UNIVERSITY
*hildebrand@dal.ca*

## References

Alexander, Michelle. (2010) *The New Jim Crow: Mass Incarceration in the Age of Colorblindness*. New York: The New Press.

Andreasen, Robin O. (2007) 'Biological Conceptions of Race'. In Mohan Matthen and Christopher Stephens (eds.), *Philosophy of Biology* (Amsterdam: Elsevier), 455–81.

Appiah, Kwame Anthony. (1986) 'The Uncompleted Argument: Du Bois and the Illusion of Race'. In Henry Louis Gates Jr. (ed.), *"Race," Writing, and Difference* (Chicago: University of Chicago Press), 21–37.

Appiah, Kwame Anthony. (1996) 'Race, Culture, Identity: Misunderstood Connections'. In Kwame Anthony Appiah and Amy Gutmann (eds.), *Color Conscious: The Political Morality of Race* (Princeton, NJ: Princeton University Press), 30–105.

Appiah, Kwame Anthony. (2005) *The Ethics of Identity*. Princeton, NJ: Princeton University Press.

Armour, Jody D. (1994) 'Race Ipsa Loquitur: Of Reasonable Racists, Intelligent Bayesians, and Involuntary Negrophobes'. *Stanford Law Review*, 46, 781–816.

Armstrong, David. (1989) *Universals: An Opinionated Introduction*. Boulder, Colo.: Westview Press.

Atenasio, David. (2020) 'The Rationality of Racial Profiling'. *Criminal Justice Ethics*, 39, 183–201.

Basu, Rima. (2019a) 'The Wrongs of Racist Beliefs'. *Philosophical Studies*, 176, 2497–515.

Basu, Rima. (2019b) 'Radical Moral Encroachment: The Moral Stakes of Racist Beliefs'. *Philosophical Issues*, 29, 9–23.

Bhogal, Harjit. (2020) 'Humeanism About Laws of Nature'. *Philosophy Compass*, 15, 1–10.

Bird, Alexander. (2018) 'The Metaphysics of Natural Kinds'. *Synthese*, 195, 1397–426.

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

Bird, Alexander, and Katherine Hawley. (2011) 'What Are Natural Kinds?' *Philosophical Perspectives*, 25, 205–21.
Blum, Lawrence. (2010) 'Racialized Groups: The Sociohistorical Consensus'. *The Monist*, 93, 289–320.
Bolinger, Renée Jorgensen. (2020) 'The Rational Impermissibility of Accepting (Some) Racial Generalizations'. *Synthese*, 197, 2415–31.
Boonin, David. (2011) *Should Race Matter?: Unusual Answers to the Usual Questions*. Cambridge: Cambridge University Press.
Bovens, Luc. (2016) 'Selection Under Uncertainty: Affirmative Action at Shortlisting Stage'. *Mind*, 125, 421–37.
Boylan, Michael. (2004) *A Just Society*. Lanham, MD, and Oxford: Rowman and Littlefield.
Boylan, Michael. (2011) 'Ethical Profiling'. *Journal of Ethics*, 15, 131–45.
Bright, Liam Kofi. (2017) 'Logical Empiricists on Race'. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 65, 9–18.
Cartwright, Nancy, and Jeremy Hardie. (2012) *Evidence-Based Policy*. Oxford: Oxford University Press.
Chetty, Raj, Nathaniel Hendren, Maggie R. Jones, and Sonya R. Porter. (2020) 'Race and Economic Opportunity in the United States: An Intergenerational Perspective'. *Quarterly Journal of Economics*, 135, 711–83.
Climenhaga, Nevin. (Forthcoming) 'Evidence and Inductive Inference'. In Maria Lasonen-Aarnio and Clayton Littlejohn (eds.), *The Routledge Handbook of the Philosophy of Evidence* (Routledge).
Colyvan, Mark, Helen M. Regan, and Scott Ferson. (2001) 'Is it a Crime to Belong to a Reference Class?' *Journal of Political Philosophy*, 9, 168–81.
Douglas, Heather. (2000) 'Inductive Risk and Values in Science'. *Philosophy of Science*, 67, 559–79.
Egan, Andy. (2011) 'Comments on Gendler's "The Epistemic Costs of Implicit Bias"'. *Philosophical Studies*, 156, 65–79.
Epstein, Brian. (2019) 'What Are Social Groups? Their Metaphysics and How to Classify Them'. *Synthese*, 196, 4899–932.
Fullinwider, Robert. (2018) 'Affirmative Action'. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (summer 2018 ed.). https://plato.stanford.edu/archives/sum2018/entries/affirmative-action/.
Gardiner, Georgi. (2018) 'Evidentialism and Moral Encroachment'. In Kevin McCain (ed.), *Believing in Accordance with the Evidence: New Essays on Evidentialism* (Cham: Synthese Library), 169–98.
Gendler, Tamar Szabó. (2011) 'On the Epistemic Costs of Implicit Bias'. *Philosophical Studies*, 156, 33–63.
Glasgow, Joshua. (2009) *A Theory of Race*. New York: Routledge.
Goodman, Nelson. (1983) *Fact, Fiction, and Forecast*. Cambridge, Mass.: Harvard University Press. First published 1955.
Hacking, Ian. (1995) 'The Looping Effects of Human Kinds'. In Dan Sperber, David Premack, and Ann J. Premack (eds.), *Causal Cognition: A Multi-Disciplinary Approach* (Oxford: Clarendon Press), 351–82.
Hacking, Ian. (1999) *The Social Construction of What?* Cambridge, Mass.: Harvard University Press.
Harcourt, Bernard. (2007) *Against Prediction*. Chicago: University of Chicago Press.
Harris, Kimberly Ann. (2017) 'Steve Biko and the Liberatory Potential of Non-Racialism and Post-Racialism'. *Critical Philosophy of Race*, 5, 223–42.
Haslanger, Sally. (1995) 'Ontology and Social Construction'. *Philosophical Topics*, 23, 95–125.
Haslanger, Sally. (2000) 'Gender and Race: (What) Are They? (What) Do We Want Them to Be?' *Noûs*, 34, 31–55.
Haslanger, Sally. (2008) 'A Social Constructionist Analysis of Race'. In Barbara Koenig, Sandra Soo-Jin Lee, and Sarah Richardson (eds.), *Revisiting Race in a Genomic Age* (New Brunswick: Rutgers University Press), 56–69.

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press

Hildebrand, Tyler. (2020) 'Non-Humean Theories of Natural Necessity'. *Philosophy Compass*, 15, 1–14.
Hildebrand, Tyler. (2023) *Laws of Nature*. Cambridge: Cambridge University Press.
Jaggar, Alison. (1987) 'Sex Inequality and Bias in Sex Difference Research'. *Canadian Journal of Philosophy*, 13 (supplement), 25–39.
Jeffers, Chike. (2013a) 'Review of *Should Race Matter? Unusual Answers to the Usual Questions*'. *Critical Philosophy of Race*, 1, 125–30.
Jeffers, Chike. (2013b) 'The Cultural Theory of Race: Yet Another Look at Du Bois's "The Conservation of Races"'. *Ethics*, 123, 403–26.
Jelenkovic, Aline, Reijo Sund, Karri Silventoinen, *et al.* (2016) 'Genetic and Environmental Influences on Height from Infancy to Early Adulthood: An Individual-Based Pooled Analysis of 45 Twin Cohorts'. *Scientific Reports*, 6, article 28496. https://doi.org/10.1038/srep28496.
Khalidi, Ali Muhammad. (2015) 'Three Kinds of Social Kinds'. *Philosophy and Phenomenological Research*, 90, 96–112.
Khazan, Olga. (2014) 'How We Get Tall'. *Atlantic*. Accessed August 15, 2017. https://www.theatlantic.com/health/archive/2014/05/how-we-get-tall/361881/?utm_source=atlfb.
Kinney, David, and Liam Kofi Bright. (2023) 'Risk Aversion and Elite-Group Ignorance'. *Philosophy and Phenomenological Research*, 106, 35–57.
Kornblith, Hilary. (1993) *Inductive Inference and Its Natural Ground*. Cambridge, Mass.: MIT Press.
Koskinen, Inkeri. (2022) 'Reactivity as a Tool in Emancipatory Activist Research'. *European Journal for Philosophy of Science*, 12, 1–15.
Langton, Rae. (2009) 'Speaker's Freedom and Maker's Knowledge'. In Rae Langton (ed.), *Sexual Solipsism: Philosophical Essays on Pornography and Objectification* (New York, Oxford University Press), 289–310.
Lever, Annabelle. (2005) 'Why Racial Profiling Is Hard to Justify: A Response to Risse and Zeckhauser'. *Philosophy & Public Affairs*, 33, 94–110.
Levin, Michael. (1992) 'Responses to Race Differences in Crime'. *Journal of Social Philosophy*, 23, 5–29.
Lewis, David. (1983) 'New Work for a Theory of Universals'. *Australasian Journal of Philosophy*, 61, 343–77.
Lewis, David. (1986) *On the Plurality of Worlds*. Oxford: Blackwell.
Mallon, Ron. (2006) '"Race": Normative, Not Metaphysical or Semantic'. *Ethics*, 116, 525–51.
Mallon, Ron. (2016) *The Construction of Human Kinds*. New York: Oxford University Press.
McPherson, Lionel K. (2015) 'Deflating "Race"'. *Journal of the American Philosophical Association*, 1, 674–93.
Moss, Sarah. (2018) *Probabilistic Knowledge*. New York: Oxford University Press.
Munton, Jessie. (2019a) 'Perceptual Skill and Social Structure'. *Philosophy and Phenomenological Research*, 99, 131–61.
Munton, Jessie. (2019b) 'Beyond Accuracy: Epistemic Flaws with Statistical Generalizations'. *Philosophical Issues*, 29, 228–40.
Nagel, Thomas. (1973) 'Equal Treatment and Compensatory Discrimination'. *Philosophy & Public Affairs*, 2, 348–63.
Pigliucci, Massimo. (2013) 'What Are We to Make of the Concept of Race?: Thoughts of a Philosopher-Scientist'. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 44, 272–77.
Rachels, James. (1978) 'What People Deserve'. In John Arthur and William Shaw (eds.), *Justice and Economic Distribution* (Englewood Cliffs, NJ: Prentice-Hall), 150–63.
Reiman, Jeffrey. (2011) 'Is Racial Profiling Just? Making Criminal Justice Policy in the Original Position'. *Journal of Ethics*, 15, 3–19.
Risse, Mathias, and Richard Zeckhauser. (2004) 'Racial Profiling'. *Philosophy and Public Affairs*, 32, 131–70.
Ritchie, Katherine. (2013) 'What Are Groups?' *Philosophical Studies*, 166, 257–72.
Ritchie, Katherine. (2015) 'The Metaphysics of Social Groups'. *Philosophy Compass*, 10, 310–21.
Ritchie, Katherine. (2019) 'Should We Use Racial and Gender Generics?' *Thought*, 8, 33–41.

Rodriguez-Pereyra, Gonzalo. (2002) *Resemblance Nominalism*. Oxford: Clarendon Press.
Rosenthal, Robert, and Lenore Jacobson. (2003) *Pygmalion in the Classroom: Teacher Expectation and Pupils' Intellectual Development*. Norwalk, CT: Crown House Publishing Company.
Schauer, Frederick. (2003) *Profiles, Probabilities, and Stereotypes*. Cambridge, Mass.: Harvard University Press.
Sider, Theodore. (2011) *Writing the Book of the World*. Oxford University Press.
Spencer, Quayshawn. (2012) 'What "Biological Racial Realism" Should Mean'. *Philosophical Studies*, 159, 181–204.
Spencer, Quayshawn. (2014) 'A Radical Solution to the Race Problem'. *Philosophy of Science*, 81, 1025–38.
Stengel, Richard. (2017) 'Why Saying "Radical Islamic Terrorism" Isn't Enough'. *The New York Times*. February 13, 2017. Accessed May 17, 2017. https://www.nytimes.com/2017/02/13/opinion/why-saying-radical-islamic-terrorism-isnt-enough.html?_r=0.
Tahko, Tuomas. (2021) *Unity of Science*. Cambridge: Cambridge University Press.
Taylor, Paul. (2004) *Race: A Philosophical Introduction*. Cambridge: Polity Press.
Thomas, Laurence. (1992) 'Statistical Badness'. *Journal of Social Philosophy*, 23, 30–41.
Thomasson, Amy. (2019) 'The Ontology of Social Groups'. *Synthese*, 196, 4829–45.
Thomsen, Frej Klem. (2011) 'The Art of the Unseen: Three Challenges for Racial Profiling'. *Journal of Ethics*, 15, 89–117.
Thomson, Judith Jarvis. (1973) 'Preferential Hiring'. *Philosophy & Public Affairs*, 2, 364–84.
Zack, Naomi. (1993) *Race and Mixed Race*. Philadelphia: Temple University Press.
Zack, Naomi. (2002) *Philosophy of Science and Race*. New York: Routledge.
Zack, Naomi. (2015) *White Privilege and Black Rights: The Injustice of U.S. Police Racial Profiling and Homicide*. New York: Rowman and Littlefield.

https://doi.org/10.1017/apa.2023.20 Published online by Cambridge University Press