

Stereotyping Patients

Katherine Puddifoot

1. Introduction

There is a burgeoning psychological literature indicating that health professionals are influenced in their clinical judgment and decision making by implicit biases. Implicit biases are automatic and unintentional associations that are made between members of particular social groups (racial, gender, socioeconomic, and so on) and certain traits (e.g., laziness, greed, athleticism, and so on) or affective responses (i.e., positive or negative affectivity).¹ When automatic stereotyping occurs through the operation of implicit bias, people are associated with particular traits in virtue of their social group membership rather than their other personal characteristics. This stereotyping has been found to influence the judgment and decision making of health professionals, leading to differential medical outcomes. The quality of care that patients receive can be determined in part by their social group membership and the associations that are made with their social group by those responsible for patient care.

The operation of implicit biases can thus bring substantial *ethical costs*, leading health professionals to treat their patients in ways that are unfair and unjust. Ethical principles of justice and fairness demand that differential outcomes are eradicated (see, e.g., Matthew 2015). They demand that people are treated equally, and given equal access to high-quality health care, without stereotypes about the social groups to which they appear to belong impacting upon the care that they receive. Ethical demands thus *prima facie* favor interventions that prevent health professionals from being responsive to the social group status of their patients in their clinical judgment and decision making.

However, it can be crucial to successful clinical judgment and decision making that health professionals are responsive to the social group status of their patients. Certain conditions are significantly more prevalent in some social groups than others. For health professionals to make correct clinical judgments, they need to reflect the prevalence of medical conditions across different groups. In addition to this, health professionals who engage in interactions with their patients that are required to obtain the information required to make correct diagnoses and treatment decisions will often discover their social group status. For these reasons, it seems that with respect to *epistemic* goals, gaining true belief, knowledge, or understanding about patients' conditions, health professionals *should* be responsive to their patients' social group status.

At face value, then, there is a conflict between ethical goals and the epistemic goals. Several philosophers have argued that people can face an ethical–epistemic dilemma with respect to implicit bias (Kelly and Roedder 2008; Gendler 2011; Egan 2011; Mugg 2013; cf. Madva 2016 and Puddifoot 2017): people who make the ethical choice to respond in an egalitarian way can suffer epistemic costs because their judgments will not reflect the distribution of traits across a population. Health professionals seem to face a dilemma of precisely this kind. They can achieve the ethical goal of responding in an egalitarian way by treating their patients the same regardless of their social group status, but they will consequently suffer epistemic costs. Alternatively, they can achieve the epistemic goal of doing the thing that is most likely to produce correct clinical judgments by reflecting their patients’ social group status in their clinical judgment and decision making, but thereby suffer ethical costs, failing to meet principles of justice and fairness.

In this article, I argue that although it is highly intuitive this is an oversimplified view. This is because there can be ethical gains from what is, *prima facie*, the best thing from an epistemic perspective. Meanwhile, and more surprisingly, there can be epistemic costs associated with what initially seems to be the best from an epistemic perspective: Even when the social group status of the patient is relevant to a judgment about her condition, recognition of her social group status can bring substantial epistemic costs.

My claim is not that health care providers should be unresponsive to the social group status of their patients, for example, having patient files with information about social status removed. We shall see that such a strategy would have serious negative consequences for precisely those who could otherwise be the object of stereotyping. Instead, I aim to show that the only way to avoid substantial epistemic as well as ethical costs is to tackle the problem of implicit bias head-on so that health care professionals can respond to the social group status of their patients without consequently making the significant errors that currently follow from stereotyping. However, tackling the problem of implicit bias in health care will be especially difficult because some of the automatic associations which produce epistemic costs can be crucial to facilitating accurate diagnostic and treatment decisions.

2. Implicit Bias and Health Care

Let us begin by exploring how ethical principles seem to demand that health professionals are not responsive to the social group status of their patients when engaging in clinical judgment and decision making. The reason for this is apparent on considering the burgeoning empirical literature on implicit bias in health care.

As implicit biases are automatic and unintentional responses, they can influence the judgment that a person makes without the person intending for them to do so. Mere awareness of the social group status of an individual can make a person

respond to them in a way that is fitting with a stereotype of the individual's social group. In other words, implicit bias can make people engage in automatic and unintentional stereotyping, where stereotyping is defined in the following way:²

Stereotyping: making a judgment about an individual that is influenced by a mental state associating members of a group, to which that individual is perceived as belonging, more strongly than members of other groups with particular traits.

It has been found that health professionals are often prone to implicit biases, associating their patients with certain traits due to their social group status, and these associations can influence the judgments they make about the patients and their conditions. Based on the empirical findings, it seems as if implicit biases influence assessments of pain, the quality of clinician–patient interactions, treatment choices, and which conditions are considered to provide a plausible explanation of a patient's condition.

Evidence that health professionals' assessments of the pain of their patients can be influenced by implicit biases is found in a study undertaken by Vani Mathur et al. (2014). These authors found that experimental participants implicitly primed with a picture of a Black face prior to considering a case study perceived and responded to the pain described in the study less than those who were implicitly primed with a White face. Meanwhile, Adam Waytz, Kelly Marie Hoffman, and Sophie Trawalter (2015) measured the extent to which White participants engaged in implicit super-humanization of Black people, that is, implicitly associating them with superhuman qualities. They found that Black people were more strongly associated with superhuman qualities than White people and, correspondingly, were less likely to be thought to require pain medication.

Another set of studies focuses on the quality of the physician–patient interaction, suggesting that it can be determined by the implicit racial bias harbored by the physician. In a study undertaken by Lisa Cooper and colleagues (2012), physicians' implicit racial bias, and their bias associating members of different races with compliance or noncompliance, were measured alongside their behaviors in physician–patient interactions. Where physicians were found to have higher levels of implicit bias associating Black people with noncompliance, the dialogue between the patient and the physician was rated by a third party as less patient-centered. That is, the patient was given less opportunity to speak. Similarly, Irene Blair et al. (2013) measured primary care clinicians' implicit bias, finding that two-thirds displayed a bias favoring Whites more than Blacks and Latinos. They then surveyed patients who had regular contact with those clinicians over a period of on average three years, asking for their ratings of the clinicians. They found that Black patients' ratings of their clinicians' patient-centeredness were predicted by the strength of the clinicians' implicit preference for Whites over Blacks (see also Penner et al. 2010; Hagiwara, Kashy, and Penner 2014).

There is strong evidence that implicit biases like these influence treatment choice. In an early study on implicit bias in health care, high levels of negative

implicit bias toward Black people predicted low levels of recommendation of thrombolytic drugs to Black patients (Green et al. 2007). In a more recent study, Jacqueline Nolan et al. (2014) explored the reasons for discrepant rates of cervical cancer screenings and follow-up care for Black women in Massachusetts in the United States. They surveyed Black women from a variety of backgrounds and found that they tended to explain the discrepant treatment by appeal to unconscious biases, with some citing a physicians' desire not to touch them as a reason why they were given inadequate care. In another study, undertaken by Diana Burgess et al. (2014), male physicians were found to be less likely to prescribe opioids to Black patients than White patients when under a condition of cognitive load, in which they were required to complete a memory exercise under time constraints while responding. It has previously been found that implicit biases are more likely to manifest under conditions of high cognitive load, time pressures, and high stress (Bertrand, Chugh, and Mullainathan, 2005; White, 2014). It is therefore highly plausible that the discrepant prescription choices were the result of the operation of implicit bias.

A final set of findings suggests that implicit biases can also determine which conditions are considered in a diagnostic process. Gordon Moskowitz, Jeff Stone, and Amanda Childs (2012) primed medical doctors with images of African American or White faces, which they were shown at speeds so fast that they were not aware that they had been shown images at all. Then they engaged in a categorization task. They were presented with words and required to categorize them as either names of diseases, names of treatments or as neither. Their reactions times were measured and compared. The doctors were significantly faster to respond to diseases stereotyped as African American (in an earlier study of some doctors' knowledge of diseases strongly associated with African Americans) when they had been shown an African American face than when they had been shown a White face. The results reveal a priming effect: exposure to faces of African Americans decreased response times to diseases stereotyped as more strongly associated with African Americans seemingly because the recent exposure to the faces primed the participants to think about the stereotypical diseases. This strongly suggests that the doctors automatically associate some diseases more strongly with African Americans than others. It also suggests that health professionals who associate particular social groups with certain conditions are likely to apply the stereotypes regardless of whether or not they intend to. On mere exposure to members of the social group, conditions that are associated with the group will be made more cognitively accessible to the health professionals, so they are more likely to automatically attribute the symptoms of the patient to the condition.³

Taken together, these results provide compelling evidence that health professionals' engagement in automatic and unintentional stereotyping can influence the judgments they make about their patients.⁴ Whether a patient is treated well or badly, and whether or not they get treatment appropriate to their needs, can depend upon the associations made by their health care provider with the social group(s)

to which the patient is perceived as belonging. There is clearly a strong case for saying that this is an unethical outcome, violating principles of justice and fairness in an area of life that is crucial to health, well-being, and even survival.⁵ Because the stereotypes leading to these negative outcomes operate automatically once a health professional is responsive to the social group status of their patients, there is reason for thinking that it will be ethically costly, that is, it will reduce the chance of health professionals acting fairly and justly, if they are responsive to the social group status of their patients in their clinical judgment and decision making.

3. Medical Conditions and Social Groups

It might seem, then, that interventions should be developed to prevent health professionals from being responsive to the social group status of their patients, so that they base their judgments wholly on the symptoms that the patient describes and the testimony that they provide about their medical history. However, this section will show that information about the social group status of patients can be highly relevant to medical judgment and decision making and a crucial determinant of whether or not a correct judgment is made. Moreover, for health professionals to gain the information that they require to make accurate clinical judgments and decisions they will often need to be placed in a situation in which they will inevitably become aware of and respond to the social group status of their patient. Therefore, substantial epistemic costs would be incurred if health professionals were placed in a situation in which they were not likely to be responsive to the social group status of their patients.

To understand the first of these points, consider how medical conditions are often unevenly distributed across social groups. For instance, hypertension, coronary heart disease, osteoarthritis, diabetes, and certain types of cancer are more commonly found among obese people than underweight, normal weight, or overweight people (Sturm 2002). People of low socioeconomic status in some countries are more vulnerable than members of other social groups to certain conditions such as tuberculosis, HIV-AIDS, and diabetes (Root 2000, cited by Sally Haslanger 2004). Meanwhile, as members of certain ethnic and racial groups, such as Black people in the United States, are statistically more likely to live in poverty, they are more likely than average to have conditions commonly found among those from low socioeconomic backgrounds (*ibid.*).

Blacks are seven times more likely to die of tuberculosis than Whites, three times more likely to die of H.I.V.-A.I.D.S. and twice as likely to die of diabetes. The diseases are biological but the racial differences are not; How is this possible? ... No mystery. Race affects income, housing, and health care, and these, in turn, affect health. Stress suppresses the immune system and being Black in the United States today is stressful. (Root 2000, S629, cited in Haslanger 2004, 11)

Because medical conditions are not evenly distributed across social groups, if a patient is a member of a particular social group, they can consequently be significantly

more likely than nongroup members to have a particular condition. Under these circumstances, both information about the social group membership of a patient and information about the prevalence of a medical condition within patient's social group is relevant background statistical information (i.e., base rate information). If members of social group A are more likely than members of other social groups to have a certain condition C, and patient P is a member of social group A, then P's social group membership is relevant to a judgment about whether they have condition C. Information about the prevalence of a medical condition can, for example, guide judgments about which diagnostic hypotheses are more probable than others. If diseases x, y, and z are more common among a patients' social group than other diseases, and a patient displays symptoms consistent with x, y, and z, the patient is more likely to have one of these diseases than some other disease, all else being equal. Health professionals cannot devote equal time and attention to all possible diagnostic hypotheses. This means that they are more likely to identify the correct diagnosis of a patient's disease if they focus on specific, probable, diagnostic hypotheses. If diseases x, y, and z are significantly more common among a social group than other diseases, then the diagnostic hypothesis that a patient of the social group has one of the diseases will be probable relative to the probability that the patient has another disease. The health professional who responds to the social group status of her patient by reflecting the information about the distribution of the diseases among the patient's social group in her diagnostic judgments can therefore be more likely to identify the correct diagnosis for the patient than a medical practitioner who does not do so. A health professional who does not reflect information about the distribution of medical conditions across different social groups in her clinical judgments might consider a large and varied range of diagnostic hypotheses. However, without a restriction on the scope of conditions considered to reflect the social group membership of a patient, the health professional is likely to fail to consider some clinical hypotheses that are highly likely given a patient's social group membership because she will be considering very many hypotheses that are less probable.

In some scenarios, information about the social group membership of a patient will be equally if not more diagnostic than information about their symptoms and medical history. Take a patient from a low socioeconomic background in a city in which there has been a high incidence of lead poisoning in areas of high deprivation due to poor housing conditions. Lead poisoning is difficult to diagnose and the symptoms of the condition can also be indicative of a number of other conditions (for example, some key symptoms are high blood pressure, difficulties with memory and concentration, headaches). A health professional could have knowledge of the symptoms and medical history of the patient without considering lead poisoning as a plausible explanation of those things. Meanwhile, knowledge of the social group status of the patient might lead her to quickly consider that the patient has lead poisoning.

To see the second point, consider how the process of gaining the information needed to make correct diagnostic and treatment decisions will frequently require direct communication between health professionals and patients and sometimes

also a physical examination by the health professional of the patient. If either of these interactions happen, then the health professional is likely to discover at least some of the social group statuses of their patients, from their appearance, accent, and so on. This means that it will not be possible to do what is required to make correct diagnostic and treatment decisions without occupying a situation in which the social status of the patient is apparent. As has been noted above, mere recognition of a person's social status can lead to an automatic response. Therefore, it will often not be possible to get the information required to make correct diagnostic and treatment decisions without automatically responding to the social status of one's patient.

In sum, health professionals who are responsive to the social group status of their patients are often significantly more likely to make correct clinical judgments. For this reason, it seems that if one is concerned with health professionals achieving *epistemic goals*, like making accurate judgments about the medical conditions of their patients, then one ought to accept that it is important for health professionals to be responsive to the social group status of patients.

4. Ethical–Epistemic Dilemma?

At this point, there appears to be good reason to believe that there is a tension between ethical and epistemic goals that must be addressed by those aiming to improve clinical practice. Ethical principles seem to dictate that health professionals' responses to patients are not determined by the patients' social group status. In contrast, the achievement of epistemic goals seems to require that health professionals be influenced by the social group status of their patients when making judgments about their condition and care. The subsequent sections show that things are not this simple. First, health professionals frequently cannot act ethically in their clinical judgment and decision making without responding differentially to patients based on their social group status. Second, responding to a patients' social group membership can bring substantial epistemic costs, so is not wholly good from an epistemic perspective.

5. How the Epistemically Best Option Can Be Ethical

This section focuses on showing that health professionals cannot act ethically without responding differentially to patients based on their social group status. In other words, the achievement of ethical goals requires what initially appeared to be the best thing from an epistemic but not an ethical perspective, that is responding to the social group status of a patient. People who adopt the role of medical practitioner take on the responsibility of identifying and responding to patients' clinical needs. Once they take on this responsibility, the ethical goal of treating people fairly demands that they do all that they can within reasonable limits to make accurate judgments about their patients' needs. If they are more likely to make correct judgments by being responsive to the social group status of their

patients, responding differently in their judgments of members of different social groups, then there can be an ethical demand on them to do so. The ethical goal of treating people fairly can therefore sometimes only be achieved via the fulfillment of the epistemic goal of making a correct judgment. Where the epistemic goal requires reflecting the social group status of a patient in their judgments, there can be an ethical demand on health professionals to be responsive to social group status.

To see this point more clearly, consider the inadequacy, in the current context, of the response made by Andy Egan (2011) to the claim that people can face a dilemma with respect to implicit bias. Egan accepts that ethical and epistemic principles clash and responds by arguing that people should be willing to make inaccurate judgments in order to treat people fairly. Egan's suggestion is in a long tradition of arguments to effect that epistemic goals can be overridden for sake of other goals (Code 1987), but it is unsatisfactory in the current case. A medical practitioner who takes the risk of making an inaccurate clinical judgment in order to be egalitarian could rightly be accused of negligence, failing to meet the responsibilities of her post. She could fail to achieve ethical goals by neglecting epistemic goals.

6. Epistemic Costs of Responsiveness to Evidence

Perhaps it is not surprising that health professionals can be required, by principles of justice and fairness, to do all that they can to ensure that they make correct clinical judgments and treatment decisions. What is likely to be more surprising is that responding to facts about a patients' social group membership can bring substantial *epistemic* costs. Moreover, these costs can be wrought even where information about social group membership is relevant to a patient's clinical diagnosis and treatment.

As we have already seen, health professionals who are sensitive to the social group status of their patients are likely to engage in implicit stereotyping. This section shows how this implicit stereotyping can bring significant *epistemic costs*. In this discussion, as in others (see, e.g., Bortolotti 2016; Puddifoot and Bortolotti 2018), x is viewed as viewed as epistemically costly if, as a consequence of x , one is less likely to achieve certain epistemic goals, including acquiring new true beliefs, retaining and using relevant information, increasing the coherence of a set of beliefs, and gaining understanding (Puddifoot and Bortolotti 2018).⁶ My claim is therefore the following: responding to the social group status of a patient can bring substantial epistemic costs. These are the epistemic costs that often follow from implicit stereotyping.

Let us begin by considering how a medical practitioner is likely to respond in a way that is influenced by implicit bias if she is sensitive to social group status of a patient. The empirical evidence cited in Section II shows that responses influenced by implicit biases often occur. Moreover, studies like the Moskowitz et al. (2012) study, also cited in Section II, strongly suggest that mere exposure

to a prime associated with a particular social group (e.g., racial group) can lead a medical practitioner to consider some diagnostic hypotheses and clinical treatments rather than others due to them making implicit associations with members of that social group.

It is not surprising that implicit biases influence clinical judgments given that the conditions of contemporary health care mean that health professionals frequently face extremely heavy workloads, which they have to complete under significant time limitations, while undergoing stress and suffering from exhaustion (Stone and Moskowitz 2011; Byrne and Tanessini 2015). Under these types of conditions, people tend to depend on implicit forms of cognition, which operate quickly and automatically (Byrne and Tanessini 2015). They can depend on associative patterns of thought rather than thinking in a controlled, deliberate manner (*ibid.*). In other words, they can automatically depend on stereotyping without having the time and opportunity to engage in other forms of reasoning and decision making.

Whether or not they intend it to be the case, then, health professionals who respond to the social group status of their patients can engage in automatic stereotyping of those patients. The following epistemic costs of responding to social group status of a patient are ways that a medical practitioner can be less likely to achieve particular epistemic goals due to the operation of these stereotypes.

A first epistemic cost is that where the stereotype that is applied associates a patient with negative characteristics, a health professional can fail to give adequate attention and credibility to the testimony of the patient. Negative stereotypes, like the stereotype that *obese people lack will power*, can lead health professionals to communicate less effectively with a patient, giving the patient less time and opportunity to explain their symptoms and generally making the patient less comfortable explaining their condition. Recall the findings discussed in Section II about the poor quality of interactions between physicians and patients when the physicians are influenced by implicit biases relating to the group(s) to which their patients belong—for example, the findings suggesting that dialogue is less patient-centered when a physician harbors a negative stereotype relating to the patient's social group. These results suggest that where health professionals are influenced by negative stereotypes, they are less likely to enter into productive dialogue with their patients, listening to the details that they can provide about their condition. They can fail to give the details that the patient provides about their condition appropriate credibility due to a stereotype they associate with the patient's social group, demonstrating what Miranda Fricker (2007) calls *testimonial injustice* (see also Carel and Kidd 2014). When people are negatively stereotyped they can also choose to smother their testimony, choosing not to provide information to people who they believe will not engage in appropriate uptake of it (Dotson 2011). Within the health care setting, this *testimonial smothering* can manifest as patients choosing not to disclose information about themselves, their symptoms, and their medical history, because they believe that the information

will be either ignored or misinterpreted by their physician who they perceive to be negatively stereotyping them. Each of these phenomena can be costly because many medical practices are fundamentally aimed at soliciting and receiving knowledge via testimony from a patient. Health professionals aim to understand this information, critically evaluate it, and give appropriate weight to it. If they engage in stereotyping, however, they are less likely to access, appropriately evaluate, or give appropriate weight to it.

A second epistemic cost of implicit stereotyping is that a medical practitioner can focus on characteristics of a patient (i.e., symptoms and aspects of medical history) that fit stereotypes associated with the patients' social group—for example, an unhealthy lifestyle in a person with low socioeconomic status—giving inadequate attention to non-stereotypical characteristics. This is because people have a general tendency to notice and attend to stereotypical characteristics more than non-stereotypical characteristics (Rothbart, Evans, and Solomon 1979; Srull, Lichtenstein, and Rothbart 1985). People also have a general tendency to build explanations that are consistent with stereotypes when other explanations would fit better with the available information (Duncan 1976; Sanbonmatsu, Akimoto, and Gibson 1994). Health professionals who are responsive to the social group status of their patients are therefore likely to explain some medical symptoms by appeal to stereotypes, for example, poor lifestyle choices in members of certain social groups, even when more adequate explanations of the information of the symptoms and medical history displayed, for example, in terms of genetic factors, are available.

It might be thought that errors of the type just described can be avoided if health professionals only focus on the *relevant* information about the social group status of their patients—the statistical information about the distribution of medical conditions across a population—and are not be influenced by stereotypes associating members of particular social groups with negative characteristics like untrustworthiness or poor lifestyle choices.⁷ It might be thought that if health professionals focused only on the relevant information then their responsiveness to the social group status of their patients would only be positive from an epistemic perspective.

There are two problems with this type of response. First of all, psychological results suggest that health professionals who are responsive to the social group status of their patients, perceiving them as members of a social group, for example, a Black patient, an obese patient, or a poor patient, will be highly likely to associate those patients with a cluster of characteristics. As noted above, due to contemporary conditions in health care, they are likely to depend on implicit stereotyping when they are responsive to social group status in this way. According to dominant theories about the psychology of stereotypes, any act of stereotyping will associate an individual with a cluster of characteristics and not one (see Puddifoot 2017). This means that if health professionals rely on stereotypes to associate members of social groups with medical conditions, they will also consequently associate those individuals with numerous other characteristics (Blair,

Ma, and Lenton 2001, cited in Moskowitz et al. 2012). Some of these associations will be epistemically costly. For example, if the stereotype that an African American patient is more likely than other patients to have sickle cell anaemia is triggered, the general stereotype of an African American patient is likely to be triggered, including an association between African American patients and uncooperativeness (Green et al. 2007). Health professionals influenced by the latter association can, for instance, fail to give appropriate weight to the testimony of African American patients.

The second problem with the suggestion that there could be only epistemic benefits if health professionals only focused on the *relevant* information about the statistical distribution of medical conditions across a population is that even stereotypes that encode this information can produce epistemic errors.

Sometimes a medical practitioner will depend on a stereotype associating a social group more strongly than other social groups with a medical condition and the association will fail to reflect reality. For example, the above-mentioned study by Moskowitz and colleagues (2012) of medical doctors in the United States found that they named thirty-six diseases as stereotypical conditions found in African Americans. These included genetically based conditions such as hypertension and sickle cell disease, but also behavior-induced conditions such as drug abuse and obesity. The experimenters found that the conditions that had a genetic basis did tend to be found more frequently in African Americans. However, the conditions induced by behavior were not found at a higher rate among African Americans. Drug abuse, for example, was found to be equally frequent among White people and African Americans in a 2003 report (Substance Abuse and Mental Health Services Administration 2004, cited in Moskowitz et al. 2012). The use of some drugs such as heroin, cocaine, stimulants, or methamphetamine was found in a report from 2007 to be higher among the White population than the African American population (Office of Applied Studies, Substance Abuse and Mental Health Services Administration 2007, cited in Moskowitz et al. 2012). This means that only some of the conditions associated with African Americans are more frequently found among the African American population. Whenever stereotypes that falsely associate medical conditions more strongly with some social groups than others are applied, health professionals will give inappropriately high levels of attention to certain conditions that they wrongly associate with a particular social group, reducing the chance that they make the correct diagnosis.

Epistemic costs can follow, however, even if a stereotype associates members of a social group more strongly than others with a particular condition *and* that condition is more prevalent in that social group than others. First of all, the application of the stereotype can lead to some diagnostic hypotheses being given undue attention and others being given insufficient attention (Moskowitz et al. 2012). As mentioned in Sections II and 3, when a stereotype is applied in the process of diagnosis it influences the hypotheses considered by the health professional because certain conditions are more likely than others to be brought to mind as potential explanations of the patient's condition (ibid). As noted in Section III, this

can be beneficial, leading the health professional to focus attention on more rather than less probable explanations of a patient's symptoms. For example, if a patient is Black, they are statistically more likely to have hypertension than members of other racial groups, and in fact highly likely to have the condition (41 percent of the Black population in the United States have hypertension compared to 27 percent of the White population) so it could be beneficial for a health professional to place hypertension high on the list of conditions that they consider when Black patients present with symptoms consistent with the condition. The stereotype will affect the "space of theoretical possibilities" that the health professional is primed to consider, which could be a positive thing if the space of possibilities that is considered both reflects the statistical distribution of conditions across different populations and is sufficiently broad. However, the same phenomenon can be costly under other conditions. Given the time pressures that health professionals operate under, the number of potential hypotheses that can be considered is severely limited. This means that they will often be susceptible to failing to formulate the correct clinical hypothesis. The application of a stereotype associating members of certain social groups with particular medical conditions can greatly increase the chance that the health professional will fail to give adequate attention to non-stereotypical clinical hypotheses because many hypotheses will be missed. Under these circumstances, health professionals are likely to fail to consider clinical hypotheses that would yield a correct diagnosis.

The application of the stereotype can also have a distorting effect on the way that a health professional perceives the symptoms of a patient. As mentioned above, stereotypes determine what information people attend to; people sometimes attend to and remember information that confirms their stereotypes but not information that disconfirms them (Rothbart et al. 1979; Srull et al. 1985). In the medical setting, this can manifest as health professionals focusing on symptoms of the patient that are indicative of conditions that they associate with the patient's social group than other social groups, failing to give adequate attention to symptoms of a patient that do not fit the stereotype. For example, a health professional might attend closely to symptoms that fit hypertension when diagnosing an African American patient but fail to attend to symptoms that do not fit hypertension, or symptoms that better fit an alternative diagnosis. Even where a stereotype reflects the statistical reality, of the high prevalence of a condition within a social group (such as hypertension among the Black population), it can have a distorting influence, reducing the chance of correct judgments being made about some patients (i.e., those with conditions that are not stereotypical).

The health professional is not likely to only attend to symptoms that are consistent with the patient having hypertension, or another stereotypical condition, if they display with symptoms that are clearly unrelated to a stereotypical condition, for example, a sore elbow. However, there will be many cases where symptoms are likely to be missed as a result of stereotyping. Perhaps the patient does not disclose information about a symptom but would provide information that would

indicate the presence of the symptom if asked the correct questions, or the presence of the symptom would be revealed if the correct tests were undertaken. Or perhaps a symptom is not easily quantifiable, such as pain, so cannot easily be conveyed to the medical practitioner. Under such circumstances, if the practitioner engages in stereotyping of social groups and conditions, psychological results suggest that they are susceptible to failing to notice or attend appropriately to the symptom if it is not stereotypical, to failing to ask the correct questions or undertake the tests necessary to reveal the presence of the symptom. Even if a stereotype that is applied reflects a genuine correlation between a condition and a social group (e.g., hypertension and the Black population), applying the stereotype can nonetheless provide an impediment to effective clinical diagnosis by preventing the practitioner from noticing and properly attending to what can be more predictive information: about non-stereotypical symptoms that provide a strong indication about which condition is present.

Even if it were the case that a health professional were to notice and attend to all relevant symptoms there could still be epistemic costs that follow from applying a stereotype about the types of conditions that members of a social group are likely to have. This is because generally when people apply stereotypes (Duncan 1976; Sagar and Schofield 1980), including automatically and unintentionally (Devine 1989), they are disposed to interpret ambiguous behavior displayed by the target of the stereotyping in a way that is fitting with the stereotype, where the behavior should be viewed as ambiguous. In the health care setting, this phenomenon can manifest as health care professionals interpreting ambiguous symptoms in ways that are consistent with stereotypical conditions being present. For example, if a Black person has a symptom that is ambiguous between a number of different conditions, but could be viewed as indicative of hypertension, the health care professional is likely to interpret the symptom as providing a higher level of support for hypertension as the correct clinical diagnosis than other conditions that the symptom is compatible with. This is epistemically costly because the correct response to ambiguous evidence is to treat it as ambiguous.

It might initially seem puzzling how it could be harmful from an epistemic perspective for health professionals to interpret symptoms that are ambiguous between different conditions in a way that is consistent with a person having a condition that is prevalent among their social group. And in some cases, a condition might be so highly prevalent in a particular population in contrast to other conditions that the presence of a symptom that might under other circumstances indicate other medical conditions provides a good indication of the presence of the stereotypical condition. For example, in a population with extremely high levels of the Ebola virus, fever-like symptoms might strongly indicate the presence of Ebola, so that it is beneficial for a medical professional to interpret fever, which would otherwise be an ambiguous symptom, as indicative of the virus. Here are two responses that can be made to these objections. First, if a condition is sufficiently prevalent in a population it is not clear that symptoms consistent

with the condition should be viewed as ambiguous. If sufficiently many people within a population have Ebola, the presence of a fever might be taken to provide unambiguous, if not inconclusive, evidence in support of the conclusion that a particular person has Ebola.

Second, if a symptom *S* is truly ambiguous between conditions *x*, *y*, and *z*, then each of conditions *x*, *y*, and *z* should be weighted as more probable than it would otherwise be due to the presence of *S*, unless there are some other good reasons for thinking that one of the conditions is not present. To see this point, it will be useful to compare it to a criminal case. There are three suspects of a crime: Johnston, Robertson, and Thompson. Some evidence is found suggesting that the criminal was wearing a red jumper. Each of the three suspects is known to have been wearing a red jumper at the time of the crime. The appropriate response to the evidence about the jumper is to weight as more probable each of the options: that Johnston is guilty, that Robertson is guilty, and Thompson is guilty. There might be independent reasons for thinking that one of the suspects is guilty, but the presence of the red jumper should not be interpreted in a way that is consistent with that suspect being guilty, if it is truly ambiguous. Instead, the evidence that each of the suspects wore a red jumper should be considered alongside, but independently of, other evidence suggesting that one of the suspects is guilty. Otherwise, the other evidence that supports the conclusion that the suspect is guilty is influential twice over, in a case of “double book-keeping”⁸ because it influences the way that the evidence about the red jumper is interpreted. Similarly, if there are three conditions that are consistent with a particular symptom, evidence that a symptom is present should lead each of the conditions to be taken equally more seriously as potential explanations of the symptoms. If there is some independent reason for thinking that one of the conditions is more likely than the others to be present, because it is prevalent within the patient’s social group, then this evidence should be considered independently but alongside the presence of the symptom. Otherwise, the fact that the condition is prevalent within the patient’s social group is influential twice over, once again in a case of “double book-keeping.”

This section has thus identified numerous epistemic costs that are associated with responding to the social group status of a patient even where information about the social group status of the patient is relevant to judgments about their condition. Health professionals who respond to information about their patients’ social group status tend to engage in stereotyping, associating the patients with a cluster of characteristics. They can consequently fail to give some patients appropriate opportunity to communicate information about their condition; they can attend closely to stereotypical features while failing to attend to non-stereotypical features; they can fail to give adequate attention to certain medical hypotheses, and so on.

The fact that some of these epistemic costs can occur as the result of an automatic association between social groups and medical conditions when the

association reflects the reality of the distribution of conditions across different social groups highlights a particularly serious problem for those who aim to increase the chance that health professionals achieve their epistemic goals. The problem is this: it will not be possible to prevent health professionals from making all of the stereotypical associations that bring epistemic costs without preventing them from making an association that it is extremely valuable for them to make.

7. In Place of the Ethical–Epistemic Dilemma

Evidence of implicit bias in medicine initially appeared to present an epistemic–ethical dilemma, but now it is clear that the situation is even more complex. The goal of treating patients in an ethical manner will often require being responsive to information about their social group status, that is doing what was initially thought only to be required in order to achieve one’s epistemic but not ethical goals. Only by being responsive to this information can health professionals discharge their duty of doing all that they can to ensure that their patients get appropriate treatment. But being responsive to this information can bring substantial epistemic costs due to the operation of stereotypes. Therefore, what initially seemed to only bring epistemic benefits—responding to the social status of patients—brings a mixture of epistemic costs and benefits.

It might seem that due to the epistemic costs associated with being responsive to the social group membership of patients, and consequently stereotyping, it would be better if health professionals were denied information about the social group membership of their patients. For instance, they could be required to make diagnostic and treatment decisions based on blinded patient files lacking details about social status. However, as discussed in Section III, health professionals often need to communicate directly with and engage in physical examination of patients to gain the information that is required to make correct diagnostic and treatment decisions. Moreover, correct judgments and decisions will often depend upon knowledge of the social group status of patients, wherever conditions are unevenly distributed across social groups.

It is worth adding to these observations that those groups that are most susceptible to being negatively stereotyped will often be those that would suffer the most if patient files were blinded, or other methods were used to prevent health professionals being aware of their patients’ social group status. Members of minority, stigmatized, and marginalized groups are most likely to be stereotyped by their health professionals, so it might seem that they would benefit the most from, for example, patient files being blind. However, the prevalence of conditions within *their* social groups are least likely to be represented in the default norms explicitly or implicitly used in diagnosis where social group status is not taken into consideration. As the probability that they will have a condition is not likely to be represented in default norms, they are most likely to be misdiagnosed if their social group status is not recognized. This is because the default norms

that do not reflect the prevalence of conditions within their group will be applied to them.

There are therefore numerous very good reasons not to conclude on the basis of the observation that there are epistemic costs associated with responding to the social group status of patients that health professionals should be denied the opportunity to respond to social group status.

It remains likely that, due to the vast number of substantial epistemic costs associated with responding to a patient's social group status that were outlined in Section VI, there are occasions when health professionals would be more likely to make correct judgment if they were unaware of the social group membership of their patients. However, it is an open empirical question how often this will be the case. And for any given clinical encounter, it will be extremely difficult to identify if it would be better for the health professional to be unaware of the patient's social status. It will be especially difficult for a health professional to make an assessment of whether they should remain unaware of a patient's social status without becoming aware of their social status. Therefore, given the importance of the interactions between patients and health care professionals that lead to acknowledgment of the patients' social status, it would be unwise to advise that practices like introducing blinded patient files should be adopted to prevent health professionals from responding to their social group status.

What should be done instead? In order to ensure that health professionals can achieve both (i) their ethical goals of treating patients justly and fairly, discharging their duty to their patients, and (ii) their epistemic goals of making correct clinical judgments and decisions, it will be necessary to reduce the extent to which stereotypes negatively influence their judgments once they are aware of a patient's social group status. Often in discussions of implicit bias and stereotyping the strategy of preventing people from being aware of the social status of group members against whom they might be biased is viewed as promising (see, e.g., defenses of making CVs anonymous). But this strategy would be detrimental in the types of cases currently under discussion.

Other strategies have been advocated to prevent the negative effects of implicit bias and stereotyping. Strategies that focus on changing the psychologies of individuals have been proposed. For example, it has been argued that considering counterstereotypical examples (e.g., strong women) (Blair et al. 2001) can change one's stereotypes, thereby changing the associations that are made when stereotypes are triggered in response to individuals. And it has been proposed that developing implementation intentions or "if-then plans" that specify how one will respond to a specific stimuli can change one's responses to individuals, controlling the association that is triggered in response to an individual (e.g., Stewart and Payne 2008; Madva 2016). For example, the implementation intention "if I see a Muslim then I will think PEACE" could alter one's responses to individual Muslims. These and other methods to combat implicit stereotyping have been found to be effective under experimental conditions (Lai et al. 2014), although their effectiveness, when administered within the experimental setting appears to

be short-lived (Lai et al. 2018), so there is reason to doubt that their effects will be significant outside the experimental setting.

Some of those who are skeptical about the effectiveness of individualistic strategies focus on structural features of society (see, e.g., Anderson 2013). For example, attempts might be made to increase integration of members of different social groups (*ibid.*). Increased integration can reduce the negative biases associated with certain social groups. It can, for instance, ensure that people encounter counterstereotypical members of other social groups consistently over time rather than for a limited amount of time within an experimental setting. The hope is that through increased integration, and other methods of social change, the stereotypes that people harbor are likely to be challenged and change accordingly.

There remains significant debate about which, if any, of these methods are effective at preventing people from automatically and unintentionally engaging in stereotyping. What the argument in Sections 1–6 shows is that, if some strategies are established to be effective at combating the negative effects of this stereotyping, this could facilitate health professionals achieving both their ethical and epistemic goals. But the argument presented so far also highlights a significant and unexpected challenge that is faced by those hoping to ameliorate medical practice by tackling implicit bias in health care. This is because the current discussion suggests that the ideal strategy to use to tackle the negative effects of stereotyping in health care would not eliminate altogether the implicit biases that have been found to manifest in health care. It would not even eliminate all of the implicit biases that bring substantial epistemic costs. Instead, it would lead people to control rather than eradicate some of these stereotypes: that is those relating social groups to medical conditions. Why should these stereotypes be controlled rather than eliminated? Because associating social groups with medical conditions, where the association reflects the distribution of conditions across social groups, can facilitate the quick and efficient selection of a correct clinical hypothesis.

Until control of this sort is taken over the automatic and unintentional stereotyping engaged in by health professionals, they are unlikely to achieve either their ethical or epistemic goals to the full because they will not be able to respond quickly and efficiently to what is important and relevant information—about their patients' social status—without suffering epistemic costs that reduce the likelihood that they will make correct diagnosis and treatment choices.

It might be that at least a partial solution to the problem of medical professionals being susceptible to stereotyping which leads to ethical and epistemic costs is relatively easy to enact. Research suggests that people are able to correct for the effects of stereotypes if they have time to do so. In a recent study, for example, people were primed with an image of a Black face and then asked to judge another image of a face according to how threatening it was (Rivers et al. 2018). When participants had a longer interval between having the stereotype of Black people as threatening activated by the prime and making the judgment of threat, they were less likely to apply the stereotype to the face that they were judging. The lengthier

interval between exposure to the prime and judgment seems to have provided the opportunity for participants to correct for the effect of the activation of the threat stereotype. What this suggests is that, if people are serious about improving the chances of health professionals achieving both (i) the ethical goal of treating patients justly and fairly, thereby discharging their duty, and (ii) their epistemic goals of making correct diagnostic and treatment decisions, then providing more time for these professionals to do their job could be an excellent place to start.

8. Conclusion

At first glance, there is an ethical–epistemic dilemma that is faced by those concerned with improving conduct in health care. There will be *ethical* costs if people are responsive to their patients' social group status because they will be susceptible to being influenced by implicit bias, engaging in stereotyping, and providing unfair treatment. However, if they are not responsive to their patients' social group status they will suffer epistemic costs associated with failing to gather and apply relevant evidence: specific evidence relating to the social group status of their patients, and evidence (e.g., about symptoms or physical signs of a condition) that could only be gathered through the kinds of interactions that would reveal the social group status of their patients.

What this article has illustrated, however, is that ethical costs follow if health professionals are unresponsive to their patients' social group status. If information about social group status is relevant to judgments about the likelihood that a person has a particular condition, and about appropriate treatments, then health care professionals can only fulfill their duty of care to their patients if they are responsive to this information. Meanwhile, there can be substantial epistemic costs associated with being responsive to patients' social group status. If one harbors implicit stereotypes relating to a patient's social group, which are triggered in the process of diagnosing and prescribing treatment for the patient, then one is likely to respond in a biased way to evidence that they might provide: failing to provide the patient with the opportunity to communicate information about their condition, failing to attend to non-stereotypical symptoms, failing to give attention to non-stereotypical medical hypotheses, and so on. Some of these epistemic costs can even follow due to the application of a stereotype even if it accurately reflects the distribution of medical conditions across social groups. What this means is that it can be an ethical choice to be responsive to a patient's social group status, and this ethical choice can also be good from an epistemic perspective, but it can also bring substantial epistemic costs.

What is the way out of this problem? It cannot prevent health professionals from being responsive to the social group status of their patients because health professionals who are not responsive to the social group status of their patients are unlikely to gather the information they need to make correct diagnoses and treatment choices. To improve both the ethical and epistemic dimensions of

clinical judgment and decision making, it will therefore be necessary to tackle the stereotyping operational in implicit bias head-on. Strategies to tackle implicit bias and stereotyping so far have been found to have some, limited effectiveness. Structural changes to society might be required to make further progress. One seemingly simple strategy that psychological studies suggest might be effective is giving health professionals more time to engage in clinical judgment and decision making—a point that seems particularly powerful with current conditions in health care in which people are often working under severe time constraints. Tackling implicit bias in health care effectively will be especially difficult because some of the stereotypes that are automatically activated in the health care context—that is those correctly associating members of social groups more strongly than others with particular conditions—can bring significant epistemic, and therefore ethical, benefits. Ideally their influence would therefore be controlled rather than eradicated. But whatever method turns out to be most effective, the current discussion shows that by effectively tackling implicit bias in health care it will be possible to reduce both ethical and epistemic costs, increasing the chance of health professionals achieving both ethical and epistemic goals.

I would to thank Lisa Bortolotti, Ema Sullivan-Bissett, Alexander Bird, Sophie Stammers, the audience at the Bristol Philosophy of Medicine seminar, and anonymous referees for helpful comments on earlier versions of this article. I acknowledge the support of the European Research Council under the Consolidator grant agreement number 616358 for a project called Pragmatic and Epistemic Role of Factually Erroneous Cognitions and Thoughts (PERFECT).

Notes

- ¹There has been a large amount of discussion about how exactly implicit biases should be characterized (see, e.g., Schwitzgebel 2010; Gendler 2011; Mandelbaum 2016; Levy 2015; Holroyd and Sweetman 2016). These debates do not have to be settled for current purposes. My claims relate to what can be concluded about the ethical and epistemic costs and benefits of implicit biases given what experimental studies say about them, and particularly and how they operate, regardless of their metaphysical status.
- ²There are alternative definitions of stereotyping according to which stereotypes always have a distorting effect on judgments (Blum 2004). However, I adopt the view of stereotyping most often found within recent social psychology, that stereotyping is any act that involves associating an individual with certain characteristics in virtue of their social group membership, regardless of the accuracy of the association. For a defense of this approach to stereotyping, see Beeghly 2015.
- ³The Moskowitz et al. (2012) study measures automatic associations rather than behavioral manifestations of bias (see note 2 for potential concerns about this), but it nonetheless strongly suggests that health professionals will be biased in their diagnoses because the order in which thoughts about conditions are manifest will determine the order in which they will consider clinical hypotheses.
- ⁴Doubts have been raised about the effectiveness of one main measure of implicit bias—the implicit association task (IAT)—for example, about whether it measures biases or recognition of attitudes or patterns found in one’s society, and about whether high levels of “bias” as measured by the IAT correlate with real-world behaviors. Some of these studies measure implicit biases using the IAT should therefore be treated with caution by those skeptical about this measure. However, the

studies discussed in this overview of the literature on implicit bias in medicine use a variety of measures, including measuring patterns of treatment under conditions in which implicit cognition tends to dominate (Burgess et al. 2014) and gathering testimony from Black female patients about their lived experience of bias (Nolan et al. 2014). The latter methodologies have not received the same criticism as the IAT. In fact, there is increased recognition of the importance of the latter source of evidence, that is, testimony of those who have experienced prejudice and discrimination (Holroyd and Puddifoot forthcoming).

⁵For further discussion of the psychological research on implicit bias in health care, see Matthew (2015). For a philosophical discussion, see Fitzgerald (2014).

⁶This list of epistemic costs is not supposed to be exhaustive.

⁷Dana Bowen Matthew (2015) seems to suggest that it is possible to make a division between being influenced by relevant information about social group status and being influenced by stereotypes.

⁸Thanks to an anonymous reviewer for the suggestion to put this point in this way.

References

- Anderson, Elizabeth. 2013. "Reply to My Critics." *Gender, Race and Philosophy: The Blog Symposia on Gender, Race and Philosophy* 9(2). Retrieved June 23, 2017, from <http://sgrp.typepad.com/sgrp/>
- Beeghly, Erin. 2015. "What Is a Stereotype? What Is Stereotyping?" *Hypatia* 30: 675–91.
- Bertrand, Marianne, Chugh, Dolly, and Mullainathan, Sendhil. 2005. "Implicit Discrimination." *American Economic Review* 95: 94–98.
- Blair, Irene V., Ma, Jennifer E., and Lenton, Alison P. 2001. "Imagining Stereotypes Away: The Moderation of Implicit Stereotypes Through Mental Imagery." *Journal of Personality and Social Psychology* 81: 828–41.
- Blair, Irene V., Steiner, John F., Fairclough, Diane L., Hanratty, Rebecca, Price, David W., Hirsh, Holen K., Wright, Leslie A et al. 2013. "Clinicians' Implicit Ethnic/Racial Bias and Perceptions of Care Among Black and Latino Patients." *Annals of Family Medicine* 11(1): 43–52.
- Blum, Lawrence. 2004. "Stereotypes and Stereotyping: A Moral Analysis." *Philosophical Papers* 33: 251–89.
- Bortolotti, Lisa. 2016. "Epistemic Benefits of Elaborated and Systematized Delusions in Schizophrenia." *British Journal of the Philosophy of Science* 67: 879–900.
- Burgess, Diana J., Phelan, Sean, Workman, Michael, Hagel, Emily, Nelson, David B., Fu, Steven S., Widome, Rachel, and van Ryn, Michelle. 2014. "The Effect of Cognitive Load and Patient Race on Physician's Decisions to Prescribe Opioids for Chronic Low Back Pain: A Randomized Trial." *Pain Medicine* 15: 965–74.
- Byrne, Aidan, and Tanesini, Alessandra. 2015. "Instilling New Habits: Addressing Implicit Bias in Healthcare Professionals." *Advances in Health Sciences Education* 20: 1255–62.
- Carel, Havi, and Kidd, Ian J. 2014. "Epistemic Injustice in Healthcare: A Philosophical Analysis." *Medicine, Health Care and Philosophy* 17: 529–40.
- Code, Loraine. 1987. *Epistemic Responsibility*. Hanover, NH: University Press of New England.
- Cooper, Lisa A., Roters, Debra L., Carson, Kathryn A., Beach, Mary Catherine, Sabin, Janice A., Grennwald, Anthony G., and Inui, Thomas S. 2012. "The Associations of Clinicians' Implicit Attitudes About Race with Medical Visit Communication and Patient Ratings of Interpersonal Care." *American Journal of Public Health* 102: 979–87.
- Devine, Patricia G. 1989. "Stereotypes and Prejudice: Their Automatic and Controlled Components." *Journal of Personality and Social Psychology* 56: 5–18.
- Dotson, Kristie. 2011. "Tracking Epistemic Violence, Tracking Patterns of Silencing." *Hypatia* 26: 236–57.
- Duncan, Birt L. 1976. "Differential Social Perception and Attribution of Intergroup Violence: Testing the Lower Limits of Stereotyping of Blacks." *Journal of Personality and Social Psychology* 34: 590–98.

- Egan, Andy. 2011. "Comments on Gendler's 'The Epistemic Costs of Implicit Bias.'" *Philosophical Studies* 156: 65–79.
- Fitzgerald, Chloë. 2014. "A Neglected Aspect of Conscience: Awareness of Implicit Attitudes." *Bioethics* 28(1): 24–32.
- Fricke, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Gendler, Tamar Szábo. 2011. "On the Epistemic Costs of Implicit Bias." *Philosophical Studies* 156: 33–63.
- Green, Alexander R., Carney, Dana R., Pallin, Daniel J., Ngo, Long H., Raymond, Kristal L., Iezzoni, Lisa I., and Banaji, Mahzarin R. 2007. "Implicit Bias among Physicians and Its Prediction of Thrombolysis Decisions for Black and White Patients." *Journal of General Intern Medicine* 22: 1231–38.
- Hagiwara, Nao, Kashy, Deborah A., and Penner, Louis A. 2014. "A Novel Analytical Strategy for Patient–Physician Communication Research: The One-With-Many Design." *Patient Education and Counseling* 95: 325–31.
- Haslanger, Sally. 2004. "Future Genders? Future Races?" *Philosophic Exchange* 34: 4–27.
- Holroyd, Jules, and Puddifoot, Katherine. Forthcoming. "Implicit Bias and Prejudice." In *The Routledge Handbook of Social Epistemology*, ed. Miranda Fricke, Peter J. Graham, David Henderson, Nikolaj Pedersen, and Jeremy Wyatt. New York: Routledge.
- Holroyd, Jules, and Sweetman, Joseph. 2016. "The Heterogeneity of Implicit Biases." In *Implicit Bias and Philosophy Volume 1: Epistemology and Metaphysics*, ed. Michael Brownstein and Jennifer Saul, 80–103. New York: Oxford University Press.
- Kelly, Daniel, and Roedder, Erica. 2008. "Racial Cognition and the Ethics of Implicit Bias." *Philosophy Compass* 3: 522–40.
- Lai, Calvin K., Marini, Maddalena, Lehr, Steven A., Cerruti, Carlo, Shin, Jiyun-Elizabeth L., Joy-Gaba, Jennifer A., Ho, Arnold K et al 2014. "Reducing Implicit Racial Preferences: I. A Comparative Investigation of 17 Interventions." *Journal of Experimental Psychology: General* 143: 1765–85.
- Lai, Calvin K., Skinner, Allison L., Cooley, Erin, Murrar, Sohad, Brauer, Markus, Devos, Thierry, Calanchini, Jimmy et al 2018. "Reducing Implicit Racial Preferences: II. Intervention Effectiveness across Time." *Journal of Experimental Psychology: General* 145: 1001–16.
- Levy, Neil. 2015. "Neither Fish Nor Fowl: Implicit Attitudes as Patchy Endorsements." *Noûs* 49: 800–23.
- Madva, Alex. 2016. "Virtue, Social Knowledge and Implicit Bias." In *Implicit Bias and Philosophy, Volume 1: Metaphysics and Epistemology*, ed. Michael Brownstein and Jennifer Saul, 191–215. Oxford: Oxford University Press.
- Mandelbaum, Eric. 2016. "Attitude, Inference, Association: On the Propositional Structure of Implicit Bias." *Noûs* 50: 629–58.
- Mathur, Vani A., Richeson, Jennifer A., Paice, Judith A., Muzyka, Michael, and Chiao, Joan Y. 2014. "Racial Bias in Pain Perception and Response: Experimental Examination of Automatic and Deliberate Processes." *Journal of Pain* 15: 476–84.
- Matthew, Dana Bowen. 2015. *Just Medicine: A Cure for Racial Inequality in American Health Care*. New York: New York University Press.
- Moskowitz, Gordon B., Stone, Jeff, and Childs, Amanda. 2012. "Implicit Stereotyping and Medical Decisions: Unconscious Stereotype Activation in Practitioners' Thoughts About African Americans." *American Journal of Public Health* 102: 996–1001.
- Mugg, Joshua. 2013. "What Are the Cognitive Costs of Racism?" *Philosophical Studies* 166: 217–29.
- Nolan, Jacqueline, Renderos, Tajan Braithwaite, Hynson, Jane, Dai, Xue, Chow, Wendy, Christie, Anita, and Mangione, Thomas W. 2014. "Barriers to Cervical Cancer Screening and Follow-Up Care among Black Women in Massachusetts." *Journal of Obstetric, Gynecologic, and Neonatal Nursing* 43: 580–88.

- Office of Applied Studies, Substance Abuse and Mental Health Services Administration. 2007. *National Survey on Drug Use and Health: Demographic and Geographic Variations in Injection Drug Use*. Washington, DC: US Department of Health and Human Services.
- Penner, Louis A., Dovidio, John F., West, Tessa V., Gaertner, Samuel L., Albrecht, Terrance L., Dailey, Rhonda K., and Markova, Tsveti. 2010. "Aversive Racism and Medical Interactions with Black Patients: A Field Study." *Journal of Experimental Social Psychology* 46: 436–40.
- Puddifoot, Katherine. 2017. "Dissolving the Epistemic/Ethical Dilemma Over Implicit Bias." *Philosophical Explorations* 20 (Supp. 1): 73–93.
- Puddifoot, Katherine, and Bortolotti, Lisa. 2018. "Epistemic Innocence and the Production of False Memory Beliefs." *Philosophical Studies* 176: 755–80. <https://doi.org/10.1007/s11098-018-1038-2>
- Rivers, Andrew M., Sherman, Jeffrey W., Rees, Heather R., Reichardt, Regina, and Klauer, Karl C. 2018. "On the roles of stereotype activation and application in diminishing implicit bias." PsyArXiv online preprint. <https://doi.org/10.31234/osf.io/7w8yz> (accessed January 2019).
- Root, Michael. 2000. "How We Divide the World." *Philosophy of Social Science* 67: 628–39.
- Rothbart, Myron, Evans, Mark, and Fulero, Solomon. 1979. "Recall for Confirming Events: Memory Processes and the Maintenance of Social Stereotypes." *Journal of Experimental Social Psychology* 15: 343–55.
- Sagar, H. Andrew, and Schofield, Janet W. 1980. "Racial and Behavioral Cues in Black and White Children's Perceptions of Ambiguously Aggressive Acts." *Journal of Personality and Social Psychology* 39: 590–98.
- Sanbonmatsu, David M., Akimoto, Sharon A., and Gibson, Bryan D. 1994. "Stereotype-Based Blocking in Social Explanation." *Personality and Social Psychology Bulletin* 20: 71–81.
- Schwitzgebel, Eric. 2010. "Acting Contrary to Our Professed Beliefs, or the Gulf Between Occurrent Judgment and Dispositional Belief." *Pacific Philosophical Quarterly* 91: 531–53.
- Srull, Thomas K., Lichtenstein, Meryl, and Rothbart, Myron. 1985. "Associative Storage and Retrieval Processes in Person Memory." *Journal of Experimental Psychology: Learning, Memory and Cognition* 11: 316–45.
- Stewart, Brandon D., and Payne, Keith B. 2008. "Bringing Automatic Stereotyping Under Control: Implementation Intentions as Efficient Means of Thought Control." *Personality and Social Psychology Bulletin* 34: 1332–45.
- Stone, Jeff, and Moskowitz, Gordon B. 2011. "Non-Conscious Bias in Medical Decision Making: What Can Be Done to Reduce It?" *Medical Education* 45: 768–76.
- Sturm, Roland. 2002. "The Effects of Obesity, Smoking, and Drinking on Medical Problems and Costs." *Health Affairs* 21(2): 245–53.
- Substance Abuse and Mental Health Services Administration. 2004. *Results from the 2003 National Survey on Drug Use and Health (NSDUH): National Findings*. Retrieved from <http://store.samhsa.gov/product/Results-from-the-2003-National-Survey-on-Drug-Use-and-Health-NSDUH-National-Findings/BKOAS05> (accessed November 1, 2016).
- Waytz, Adam, Hoffman, Kelly Marie, and Trawalter, Sophie. 2015. "A Superhumanization Bias in Whites' Perceptions of Blacks." *Social Psychological and Personality Science* 6: 352–59.
- White, Augustus, III. 2014. "Some Advice for Minorities and Women on the Receiving End of Health-Care Disparities." *Journal of Racial and Ethnic Health Disparities* 1: 61–66.