



This is a repository copy of *Responsibility for implicit bias*.

White Rose Research Online URL for this paper:  
<http://eprints.whiterose.ac.uk/111327/>

Version: Accepted Version

---

**Article:**

Holroyd, J.D., Stafford, T. and Scaife, R. (2017) Responsibility for implicit bias. *Philosophy Compass*.

<https://doi.org/10.1111/phc3.12410>

---

This is the peer reviewed version of the following article: Holroyd, J., Scaife, R. and Stafford, T. Responsibility for implicit bias. *Philosophy Compass*. 2017;12:e12410., which has been published in final form at <https://doi.org/10.1111/phc3.12410>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Self-Archiving."

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

## **Responsibility for Implicit Bias<sup>1</sup>**

Research programs in empirical psychology from the past two decades have revealed that implicit biases are pervasive. Implicit biases are typically characterised as automatic associations, of which we may not be aware, that are difficult to control, and may conflict with our professed beliefs and values.<sup>2</sup> Whilst implicit processes are pervasive, unavoidable and often useful aspects of our cognitions, they may also lead us into error. The most problematic forms of implicit cognition are those which target social groups, encoding stereotypes or reflecting prejudicial evaluative hierarchies. Despite intentions to the contrary, these biases can influence our behaviours and judgements, contributing to patterns of discriminatory behaviour. Here is a paradigmatic example in which implicit biases influence behaviour and judgment:

Professor P is engaged in anti-racist and feminist activities. She explicitly avows egalitarian values and strives for fair treatment. She believes her students deserve to be treated with equal consideration, and hasn't reflected much on her behaviour, since it seems to her there is no reason for concern. Accordingly, she believes that she adheres to her egalitarian goals, which is what she wants and endorses.

However, Professor P harbours implicit biases of the sort found by empirical psychology: associations between whites and intellect (Amodio & Devine 2006), and negative evaluative associations with black people (Olson & Fazio 2006). These implicit cognitions influence her behaviour towards her students – again, in ways we might predict given the findings of empirical psychology:

---

<sup>1</sup> This paper was produced as part of the Leverhulme Trust Project Grant, on Bias and Blame (RPG-2013-326). We are grateful for their support.

<sup>2</sup> Not all parties to the debate accept this characterisation. See e.g. Mandlebaum 2014, Levy 2015 for the claim that implicit biases likely have propositional structure.

her evaluations of the equally good work of black students is slightly less glowing than that of white students (Wood et al 2009); her interactions with black students are beset by micro-behaviours that indicate greater tension and discomfort, and as a result she is marginally less warm and less patient with them than with her white students (Dovidio et al 2002).

Professor P is not unique – but wholly typical; her colleagues show a similar pattern of evaluations and behavioural interactions, which subtly discriminate against black students. The cumulative effects of this are manifest: fewer black students occupy the top percentiles of classes in her department, fewer pursue or are encouraged to pursue further study in her subject, the student body in her discipline, at the levels at which it is selected for, are disproportionately white.

In this scenario, implicit biases are complicit in racially discriminatory behaviour, both in individual interactions, and as part of a pattern of behaviours that contribute, alongside other exclusionary factors, to sustaining the under-representation of black students in higher education. Note that higher education is not unique; similar scenarios may be found in many domains of education and employment. Implicit racial biases have been found in studies that examine biases in healthcare contexts (Green et al 2007), housing (Turner 2007; Kang 2014), the criminal justice system (Eberhardt et al 2004), financial decisions (Tetlock et al 2000, see Madva 2016a for instructive discussion), and in everyday interpersonal interactions (Dovidio et al 2002). Moreover, other implicit biases target gender, sexuality, age or disability (Valian 1999, Cooley et al 2014, Levy & Banaji 2002; Wilson & Scior 2014). It is very likely that any one individual harbours and is influenced by some sort of implicit

bias; the findings mean that we cannot rule out that, for each of us, we are influenced in our judgements or behaviours in discriminatory ways.<sup>3</sup>

Such patterns of discrimination are obviously wrong and unjust. But in remedying such wrongs, one question to be addressed concerns responsibility for implicit bias. Yet, unlike some paradigmatic forms of wrongdoing, discriminatory behaviour resulting from implicit bias is often unintentional, unendorsed, perpetrated without awareness; and the harms may be particularly damaging because they are cumulative and collectively perpetrated. Similarly, the cognitions involved in the biases are ones from which agents may be alienated, since they conflict with endorsed values; the agent may not be aware of their presence in her cognition; nor aware of their role in behaviour. So, what are we to make of questions of responsibility, for the cognitions themselves, or for actions influenced by them? In this article, we outline some of the main lines of recent philosophical thought which address questions of responsibility for implicit bias. Our main focus, in keeping with much of the literature, is on responsibility for actions influenced by such biases. These debates have import both for how we think about and respond to implicit bias; and for our understandings of responsibility. We focus on 1) the kind of responsibility at issue; 2) revisionist vs non-revisionist conceptions of responsibility as applied to implicit bias; 3) individual, institutional and collective responsibility for implicit bias. We close with a practical example, of institutional responsibility and implicit bias training, that animates some of these issues.

#### 1) What sense of responsibility?

---

<sup>3</sup> For debate over how likely this is, see Jost 2009, Greenwald et al. 2009; Oswald et 2013.

In remedying wrongs, we might be interested both in causal responsibility and moral responsibility; that is, we want to know both who or what brought about the wrong, and who is to blame for it. Notably, some have argued that attention to individuals and their implicit biases is inapt in addressing these patterns of discrimination, since the causal responsibility for such inequalities and discriminatory patterns is fundamentally structural (Haslanger 2015, Anderson 2012). Accordingly, remedies would do better to target these fundamental, structural causes.

On the other hand, others have maintained that implicit biases, and crucially, actions that manifest them, are causally responsible for some of these harms and inequalities (Saul 2013a) – and, attention should be paid to understanding the role of implicit bias in sustaining these harms and attempting to mitigate their influence.

However, in drawing attention to the role of implicit biases in sustaining harms and exclusion, it has been contentious whether individuals are always morally responsible for these harms. Saul has suggested not, writing that we should

abandon the view that all biases against stigmatised groups are blameworthy ... [because a] person should not be blamed for an implicit bias that they are completely unaware of, which results solely from the fact that they live in a sexist culture. (2013a: 55, emphasis in original)

Here, Saul focuses on whether individuals are responsible merely for having implicit biases. However, this question is distinct from the question of whether one is responsible for manifesting implicit bias in one's actions, and one may be blameworthy for one's actions even if one is not blameworthy for the presence of the implicit association in one's cognitive make up (Holroyd, 2012). In any case, Saul

suggests that once individuals are aware that they have, and may be influenced by, implicit biases, they are responsible for doing something about it:

They may, however, be blamed if they fail to act properly on the knowledge that they are likely to be biased—e.g., by investigating and implementing remedies to deal with their biases (ibid.)

Saul's remarks here exemplify both the backward looking concern with appraising individuals for their cognitions and behaviours; and attention to the forward-looking issues concerned with taking responsibility for correcting harms and mitigating costs (of course, failure to meet these forward-looking obligations may render one blameworthy, in the backward looking sense, for that failure). In the wider debate, theorists have teased apart the forward-looking and backward-looking concerns into (inter alia) the following different questions, each of which might legitimately be regarded as concerning a sense in which the agent is responsible:

- i) does the attitude or act reflect badly (or well) on the agent; is there a fault (or indeed credit) that can be attributed to her?
- ii) should the agent be regarded as blameworthy for the fault she has, or has demonstrated in her action; should she bear some cost or burden (in the form of sanction, or blame) for this?

These two questions are backward looking, and contrast with a third concern, namely,

- iii) what forward looking obligations do individuals have for dealing with the fault, or problematic behaviour?

Those familiar with the responsibility literature will be aware that the two backward-looking concerns, above, have sometimes been characterised as questions of attributability (i) and accountability (ii) (cf. Watson, 1996, Smith 2012). However, these two terms have (confusingly) garnered somewhat different usage in the literature on responsibility for implicit bias. Zheng has labelled the backward looking senses (i and ii) as together questions of attributability; and characterised as accountability the forward looking issue of how remedial obligations are distributed (this way of using the two notions, Zheng notes, has more in common with the way Scanlon uses the terms).

All three questions are activated in thinking about responsibility for implicit bias. To avoid terminological confusion, in the following discussion, we make reference to the particular question at issue: does it reflect on the agent; is she blameworthy for it (the backward-looking questions); do obligations to address bias attach to an agent (the forward-looking concerns)?

Zheng argues that we arrive at different senses of responsibility via different routes. The backward-looking questions concern what reflects well or badly on the agent, such that they are subject to ‘appraisive responses’ such as praise or blame (or perhaps other evaluations of character). Blameworthiness, thus construed, is a backward looking matter. For example, to find Professor P responsible in this sense would be to suppose that her discriminatory evaluations and behaviours flow from the relevant features of her agency – some metaphysical or psychological facts – which allow those actions to reflect on her as an agent, and enable us to perhaps find her blameworthy for those harmful acts she perpetrates.

In contrast, the forward-looking concerns (what Zheng calls accountability) track not deep facts about the agent and what reflects on her, but rather social facts

about the distribution of benefits and burdens when something has gone wrong: who should step in to deal with the costs of wrongdoing, irrespective of whether they were at fault. We might hold Professor P to be charged with acting to remedy the costs of the harms she is implicated in, irrespective of whether we think those harms are ones that reflect on her, or for which she is blameworthy.

In some cases of implicit bias, it may be clear that the biases warrant backward-looking, appraisive responses: for example, in cases where the agent has cultivated the biases, this would reflect badly on her, and she may be blameworthy for her biases. Likewise, in cases in which the agent endorses or would endorse the discriminatory behaviour (Zheng 2016, Holroyd 2016), the behaviour appears to reflect a fault on her part; moreover judgements of blame and expressions of blame – may also be engaged.

However, in other cases – which meet the paradigm illustrated by Professor P – things may not be so clear. Various authors have suggested that there are reasons for which the discriminatory behaviour may not reflect badly on the agent, and certainly not in a way that warrants blame (see e.g. Saul 2013a, quoted above; Levy 2012, 2014 and section 2 below), since the behaviour is not something the agent is aware of, nor under the agent's control. On the other hand, Brownstein sets out the case for taking actions that express implicit bias to be reflective of who the agent is, insofar as they reflect the agent's cares (which may not form a coherent set, and hence may conflict with her other cares). What the agent cares about, Brownstein suggests, is reflected in multitrack patterns of behaviour that reveal the agent's dispositions (2015: 12). As such actions that express implicit bias can be attributable to the agent, and may ground moral appraisals of her.



In any case, one might hold such agents responsible in the in the sense of having forward looking obligations to deal with the problem, irrespective of whether bias, or actions influenced by it, reflects badly on them, or is something for which they can be blamed (Zheng 2016). Withholding backward-looking appraisive judgements and maintaining that the agent is responsible only in the forward-looking sense of having remedial obligations may seem appealing for three reasons. First, one may have doubts about appraising an agent, yet insist that it is important that someone is held responsible nonetheless (though see Holroyd 2012, Brownstein 2015, and the discussion in section 2 below, for arguments addressing different versions of such doubts). As Zheng puts it, it is important that we are held to forward-looking responsibilities ‘because it is appropriate for us to clean up after our own actions when a mess has been made’, even if the mess doesn’t reflect badly on us, or warrant blame (2016: 74). When someone has been harmed by implicit bias, she is owed ‘compensation, apology, and redress’ irrespective of whether anyone is found to be at fault or blameworthy (2016: 74). (Note that it remains for proponents of this strategy to delineate senses of apology and redress that can play this function, whilst being distinct from those appraisive responses characteristic of backward-looking responsibility judgements.) Second, it may be difficult to identify, with sufficient certainty to warrant blame, whether implicit bias has played a role in discriminatory behaviour in any isolated case.

A third reason for focusing on responsibility in the forward-looking sense concerns worries about the engagement of the appraisive responses. Saul worries that blame might produce backlash or hostility. This is a common concern: Zheng (80) marshals evidence (Legault et al 2011) which suggests that high threat confrontations can produce defensiveness and greater bias (implicit and explicit); Vargas (2017)

raises worries about securing the needed buy-in for norms against bias, if agents experience backlash in response to blame that is not accepted or well-received.

The appeal of a focus on holding people responsible in the exclusively forward-looking sense, for actions influenced by implicit bias, rests on a) the extent to which we are convinced that agents are not (or cannot be identified as) responsible in the backward-looking senses (more on this in section 2 below); and b) the plausibility of concerns about negative responses to appraisive evaluations such as blame.

On this latter point, it is worth noting that appraisive responses may include a wide range of moral evaluations that need not be characterised by ‘high threat’ features (such as high emotion, or challenge to the agent’s self-conception or esteem). For a start, appraisals that focus on whether biases reflect badly on the agent (sense (i) above) may not involve blame at all. Rather, they might invoke an evaluative judgement about the agent and her character - she is cruel, or she is racist - without taking a stance on whether this is her fault. This is the sort of ‘aretic appraisal’ (Watson 1996) that has been taken to be distinct from judgements of blameworthiness, or expressions of blame. But even blaming responses may not be ‘high threat’. Instead, preliminary empirical work by Malle et al (2012) suggests that individuals are more likely to see responses as forms of blaming when they are thoughtful, calm and delivered without high emotion (in contrast, see the conception of blame outlined and critiqued by Pickard 2013, which supposes blame has an emotional ‘sting’). Finally, some initial work on reactions to blame for implicit racial bias has found that blaming responses did not increase implicit biases, and in fact significantly increased individuals’ explicit intentions to take action to address bias. This effect on explicit intentions was long lasting, persisting in a 6 month retest, in which participants continued to report stronger motivations to tackle discrimination

(see Scaife et al ms.).<sup>4</sup> In short, no ‘backlash effect’ was observed. This result may seem surprising at first. But note that it is less so if worries about blame for implicit bias rest on a mistaken paradigm of blaming communications: one which supposes that individuals are being told they are bad and wrong. Instead, blame for implicit bias might be better modelled – at least in some cases – as a communication in which individuals are alerted to the fact that they have violated a norm that they already (rightly) endorse (see Scaife et al ms. for more discussion). We might expect reactions to these different sorts of communication to elicit rather different responses. It seems that considerable work is still to be done on understanding how to model different sorts of moral response, and on empirically evaluating how they are received and reacted to.

If responses such as blame are effective in influencing individuals’ attitudes towards avoiding implicit bias, then there may be reason to deploy such responses irrespective of whether individuals in fact deserve such responses. But absent desert, the justificatory burden for blame is certainly greater. So, much hinges on whether agents are in fact blameworthy for implicit bias or implicitly biased actions. This depends on how we conceive of the conditions for this kind of responsibility, to which we now turn.

## 2) Folk vs Revisionist conceptions of responsibility

Many philosophical conceptions of responsibility aim to articulate, perhaps with some refinements, the conditions of responsibility embedded within and deployed by folk

---

<sup>4</sup> Note that these findings address one objection to communicating blame; but the issue of blameworthiness and whether one should communicate blame may be thought to come apart (though for some consequentialists, to be blameworthy just is for there to be sufficient reason to blame). There might be reasons (backlash) to withhold blame even if the agent is blameworthy; we can imagine scenarios in which one should blame (if the gains are great) even if the agent is not blameworthy.

conceptions of responsibility. Whilst there is of course much long standing philosophical disagreement over how to articulate and unpack this conception, many agree on the centrality of the following notions in determining moral responsibility, in the sense of blameworthiness, for attitudes and actions based upon them: awareness or knowledge conditions; and control conditions (which authors have spelled out in terms of reasons-responsiveness or evidence-sensitivity conditions that have featured heavily in the literature (Fischer & Ravizza 2000; Levy 2014, forthcoming)). In a paradigm case of wrong-doing for which the agent is morally responsible, the agent's attitudes are under her control in that they are suitably responsive to reasons and to evidence; she is aware of what she is doing and of the consequences of her actions; and she is able to ensure her actions reflect her endorsed values.

It is clear enough that the discriminatory actions influenced by implicit bias – such as the evaluations and interactions of Professor P – do not meet this paradigm. Professor P's implicit biases are not sensitive to her other (evidentially supported) attitudes, beliefs and values, but conflict with them; she is unaware that she discriminates; and she does so automatically, without the reflective or deliberative control we are sometimes able to exercise. This departure from the paradigm cases of moral responsibility has led some authors to claim that individuals are not responsible for actions influenced by implicit bias: such actions cannot reflect badly on them, nor are they always blameworthy for them (see Saul 2013, above).

Levy offers extensive arguments for the claim that individuals such as Professor P are not blameworthy for implicitly biased actions, because her actions are not actions of which she is appropriately consciously aware;<sup>5</sup> nor is she able to

---

<sup>5</sup> See Cameron & Payne 2010 for studies that suggest that Levy's claims have support in the folk conception of moral responsibility. Terminologically, it is important to note that Levy is

exercise reflective control over them; nor are the attitudes integrated with (or responsive to) her endorsed evaluative attitudes (Levy 2012, 2015, ms). Accordingly, Levy claims, they cannot be reflective of who the agent is, and so it cannot be appropriate to respond with certain burdens (sanction, blame, punishment) (ms.: 4).

But of course, it is debatable whether the notions of awareness, or control, deployed in such arguments are the relevant ones for blameworthiness (see Holroyd 2012), and some lines of debate have engaged these issues.

#### a) Awareness

Consider first the awareness requirement. The idea that we are responsible (in the sense of blameworthy for) only for those attitudes or actions that fall within, or are guided by, conscious awareness faces sustained critical challenge from Sher (2009). Such a condition would implausibly render us exculpated from (inter alia) any instance of absent-mindedness, inattention, or forgetting (cf Smith 2005). We should instead, Sher suggests, start from the thought that “when someone acts wrongly or foolishly, the question on which his responsibility depends is not whether he is aware that his act is wrong or foolish, but rather whether he should be” (Sher, 2009:20).

In the context of implicit bias, various senses of awareness are in circulation (taxonomised Holroyd 2014): agents may have (or lack) introspective awareness of their implicit biases; individuals may have (or lack) inferential awareness of their proneness to implicit biases – that is, an awareness based on inferences drawn from empirical studies to our own propensities. Third, individuals may have (or lack) observational awareness – that is, awareness based on the observation of their own

---

concerned with blameworthiness, even though he frames his question as concerning whether biases are ‘attributable’ to the agent (ms. p.4).

behaviours, in one off cases or in patterns of action.<sup>6</sup> Relatedly, individuals may be aware (or not) of the moral significance of their action (Levy 2014). Combined with Sher's insight, Holroyd argues that responsibility will depend not on whether individuals in fact have or lack any such sense of awareness, but rather on whether we should have any of these kinds of awareness; and in the instance in which we do not, whether any such failure is culpable.

For instance, such a strategy is deployed by Washington & Kelly (2016), who argue that agents in certain positions of power or authority – those on hiring panels, say – should have awareness of their own likely implicit biases. That is, given their role responsibilities and the impact of their decisions, they should be aware of the relevant research on implicit bias, and be able to make inferences about the ways in which their own judgements may be biased – and, crucially, the steps they may take to mitigate this. The issue is then whether such individuals are culpable when they lack such awareness.

Washington & Kelly take this to support 'externalist' conditions on responsibility, such that an individual's failing to be aware of what she should be may or may not be culpable, contingent upon her epistemic environment. In a context in which information – from which inferences can be drawn – is not readily available, the failure of awareness is not (or at least, is less) culpable. But in an epistemic environment such as ours, in which such information is - increasingly - more readily available, we are culpable for lacking awareness of the information of which we should, and can reasonably be expected to be, aware.

---

<sup>6</sup> Note, moreover, that it is a mistake to suppose a failure to report implicit biases (on self-report measures) is equivalent to lacking awareness in any one of these senses (since there may be various reasons for failures of self-report). See Stafford 2014 for discussion.

Holroyd (2014) argues that in fact, we need not demand even this inferential awareness; agents may have observational awareness even in epistemic environments in which research on implicit bias is not ‘common knowledge’. This can be gained by careful reflection on one’s own behaviour and its role in sustaining patterns of systemic injustice. Whilst the relevant observations may be more easily made when information about implicit bias is prevalent, such information is not necessary. Indeed, prior to the advent of research on implicit bias, plenty of testimonial evidence from individuals targeted by implicit racial biases could have prompted such reflections and observations (see Holroyd & Puddifoot forthcoming). And, insofar as these observations are blocked by motivated ignorance - avoidance or ignoring (perhaps subconsciously) of evidence that serves one’s goals - or self-deception, or excessive weight given to misleading introspective evidence, a failure to have this sort of awareness is culpable to some degree.

In sum, on the folk conception of responsibility, some awareness conditions may be required for individuals to be blameworthy for their wrongful actions. But it is far from clear that simply lacking conscious awareness of facts about implicit bias, or the fact that one is discriminating, is exculpatory. Rather, we might expect that individuals – at least in some epistemic contexts – have, and should have, certain kinds of awareness about their propensity to harbour and manifest implicit bias.

#### b) Control

Contention has also arisen over the kind of control necessary to ground blameworthiness for actions influenced by implicit biases. As Saul has argued, individuals lack direct control over the activation and influence of implicit bias. We can’t simply will that biases don’t affect us, and make it so. Indeed, studies have

indicated that trying to suppress implicit bias has a rebound effect (Follenfant & Ric 2010).

However, many aspects of our cognition and action for which we are held responsible are not under our direct control. This suggests that other forms of control are relevant to moral responsibility. Two questions then arise: what forms of control are necessary for moral responsibility? And do we have these form of control with respect to implicitly biased actions?

Levy (2014, 2015) has argued that agent's attitudes must be under a certain kind of rational control – they must be inferentially sensitive to, and integrate with, the other evidence-sensitive attitudes of the agent. Drawing on evidence that implicit attitudes are not inferentially sensitive in the appropriate way, he argues that actions influenced by implicit biases cannot be reflective of the agent in a way that renders her liable to blame (Levy 2015 pp.812-816). Such implicit attitudes cannot, Levy argues, be properly integrated into the agent.

It is worth noting, however, that any such lack of inferential sensitivity does not necessitate a lack of integration. Implicit biases may accord with and thereby reflect the agent's values, even if not under rational control (Holroyd 2016). The paradigm case of this would be an explicit racist who nonetheless has implicit cognitions and automatic responses, which contribute to (and function in the service of) his morally repugnant goals.

In any case, rational control may not be necessary for moral responsibility: other sorts of control may suffice in order for agents to meet any control conditions for moral evaluation. For example, Snow (2006) identifies a sort of 'intervention control' that agents may exercise over automatic processes (in the context of discussing automaticity, rather than implicit biases per se – see Holroyd & Kelly



(2016) for discussion of intervention control in the context of implicit biases). Whilst the processes themselves are not rationally guided, but are automatized and habituated, agents may intervene to halt or redirect these processes. Such control may be possible in the case of implicit biases. For example when on a hiring panel, Professor P may intervene to halt the automatic influence of implicit biases, by, say, referring to a checklist of previously agreed upon criteria, as a strategy to try to overcome any tendency to bias (cf WISELI 2012, p.6).

Other forms of indirect control may also be deployed, such as re-training or conditioning one's own cognitions via techniques such as 'implementation intentions'. This technique involves conditioning certain counter-stereotypical automatic thoughts in response to environmental cues: for example, automatizing the thought 'safe', as prompted by any black person, in order to try to overcome associations between blackness and danger or hostility (Stewart & Payne 2008). Agents may add to these cognitive props environmental props also – reshaping their social environment in a way that impacts on their cognitions in desirable ways. For example, some studies have suggested that the presence of pictures depicting counter-stereotypical exemplars reduces negative biases (Blair 2002).

These modes of indirect control – manipulating one's environment or one's cognitions in order to secure desirable patterns of thought and behaviour – have been identified by Holroyd & Kelly (2016) as forms of 'ecological control'. They argue that this kind of control is in fact mundane, oft-deployed (e.g. consider organising one's office to ward off procrastinatory tendencies), and sufficient for meeting control conditions for moral responsibility.

Whilst the kinds of ecological control so far outlined involve intentionally deploying strategies to exert control, Holroyd & Kelly argue that this is not necessary.

One form of ecological control outlined is that of unintentional indirect control: an example of this is the sort of control demonstrated by participants in Moskowitz & Li's 2011 studies, whereby agents with strong egalitarian commitments appeared to be better able to block the activation of implicit biases. This was not explicitly intended by those agents, but seemed to be an automatic function of their strongly held egalitarian commitments. Since agents who lacked such commitments were less effective at blocking the activation of implicit stereotypes, Holroyd & Kelly argue, it is to the agent's credit and attributable to her, when her values better enable the pursuit of her goals.

Whether agents who are influenced by implicit biases meet control conditions for moral responsibility, then, will depend on what notion of control one teases out of the folk conception, how defensible this notion of control is, and the extent to which it is exercised with respect to implicit biases. Note that in teasing out these notions, there is some refinement of the ideas found in the folk conception: in the articulation of the idea of 'ecological control' for example. But authors proposing these notions take such ideas to be latent in, or extensions of, the folk concept of moral responsibility. This contrasts with the strategies deployed by theorists who are explicitly revisionist about moral responsibility.

### c) Revisionism

Some authors have suggested that thinking about blameworthiness for implicit biases should motivate revisions to this understanding of the concept of moral responsibility, whereby it is acknowledged explicitly that such revisions involve a departure from seemingly intuitive thought about what it is for attitudes and actions to reflect on us, and render us liable to blame.

Two authors who propose such revisions are Glasgow (2016) and Faucher (2016). Each is motivated by the idea that we are morally responsible for actions influenced by implicit biases – but that this can only be adequately captured by revising our analyses of moral responsibility. For example, Glasgow suggests that whilst common sense suggests that the conditions for moral responsibility are invariant – they remain constant across agents and contexts – his reflections on implicit bias lead to the revisionary conclusion that they are variant. In particular, the conditions vary depending on the content of the attitude held and the action performed on the basis of it (2016: 48) (later, Glasgow suggests that the content is a proxy for the kind or degree of harm caused (2016: 56)). On this view, the content of the attitude determines whether an agent’s alienation from her attitude (or action performed on that basis) suffices to exculpate. In the case of Frankfurt’s alienated drug-user, alienation seems to suffice. But in cases where the attitude involves a relational harm – the attitudes that drive infidelity, or prejudicial attitudes involved in implicit biases – the fact that the agent is alienated from her attitude does not serve to exculpate.

Faucher (2016) also argues for a revisionary form of variantism about moral responsibility. On his view, the dimension of variation concerns the conditions that are applied depending on one’s position in relation to the harm. If one is – or has been – a victim of discrimination, the conditions one deploys in evaluating moral responsibility will not make reference to explicit or conscious intentions, he argues. Harm has been perpetrated and it matters not, from this point of view, whether it was consciously intended. However, from the point of view of perpetrators, conscious intention does matter.

One might think that this difference rests on mistake or disingenuousness on the part of the perpetrators. But if it is genuine and withstands reflective scrutiny, then our conditions for moral responsibility may turn out to be variant – and an invariantist articulation of the concept of moral responsibility should be revised accordingly.

Faucher also suggests other revisions; first, to the notion of control necessary for responsibility. As discussed above, forms of indirect control (Faucher calls those which involve retraining of one's cognitions a kind of 'bottom-up' control) may be sufficient for moral responsibility. And notions of responsibility that make reference to the 'real self' - the part of the self that reveals where the agent stands – may have to account for the agent's real self as including some unendorsed implicit attitudes, as well as her endorsed explicit attitudes.

An explicitly revisionist model of responsibility has been developed and defended by Manuel Vargas in recent years (2005, 2009, 2013). His recent work has applied this model to implicit biases. Vargas (2017) sees responsibility judgements as part of a practice that can be justified if it serves certain forward-looking goals of cultivating moral agency. On this view, to be a responsible agent is to stand in certain social relations (rather than to meet certain metaphysical conditions). Thus construed, it is crucial that the agent's context supports moral agency in relation to a certain set of considerations. Vargas worries that our current environment (the 'moral ecology') does not yet provide the right sort of support. We are perhaps not yet suitably cognisant of, or sensitive to considerations of implicit bias, to make it appropriate – or fruitful – to hold each other responsible or liable to blame, on his view. But, as our moral ecology evolves, it may be appropriate to do so. This line of argument has resonance with Washington & Kelly's (2016) externalist conditions for moral responsibility, whereby what can reasonably be expected of agents may shift as

epistemic environments change. Note also that this view about the current impropriety of blame shares much with the concerns about the defensiveness and hostility that such appraisive responses may invoke. Yet, as we have seen, it is not at all clear that we do find these responses – it may instead be that these interpersonal interactions of blame are part of the tools that shape the moral ecology and increase sensitivity to, and motivation to combat, implicit biases (Holroyd 2012, Scaife et al ms.).

These moral responses are premised chiefly on the idea that individuals may be held responsible for implicit biases. But this is not the only option.

### 3) Individual, institutional and collective responsibility

So far, the focus has been on holding individuals morally responsible for particular discriminatory actions or judgements made under the influence of implicit bias. But as noted at the outset, the problem is not simply one of particular discriminatory behaviours, but also of their role in sustaining patterns of marginalisation, exclusion, and hierarchy. Insofar as these social structures are collectively caused, sustained, and stand in need of collective remedy, it appears that an approach that requires collective responsibility is also needed.

Of course, questions remain about who is collectively responsible for what, in both the backward-looking (evaluation of character; or identification of who is to blame) and forward-looking (who bears remedial costs) senses. One might maintain that all of us are to blame for, or have a role in remedying, the harms that result from implicit biases. Indeed, as Jacobson has noted (2016, 174), agents need not have done anything particularly wrong in order to have been complicit in sustaining patterns of exclusion. This complicity may justify being held accountable for remedying the harms. Or one might maintain that responsibility and blameworthiness falls

collectively, but proportionally, on those groups who have greater role in perpetrating harm (cf. Washington & Kelly's remarks about people in 'gatekeeper' positions having greater responsibilities). Likewise, in terms of forward-looking responsibility, Sie and Van Voorst Vader-Bours suggest that 'it seems reasonable to expect those in positions of power to take more responsibility for change' (2016: 107). Moreover, if it is right that remedying implicit biases is most effectively done by making broader structural changes (Saul 2013b, Haslanger 2015, Anderson 2012), then certainly institutional and collective action will be needed to enact this. But as Madva has persuasively argued (2016b) attention to individual or collective endeavours should not be seen as competing; rather, the complex interplay between individual and more collective or structural interventions needs to be recognised. For example, institutional change requires individual buy-in and motivation to instigate change. We close with an example of how both institutional and individual responsibility need to have a role in practical efforts to remedy implicit bias.

Many institutions now offer 'implicit bias training' – seminars that aim to make individuals within an institution aware of the possibility of implicit biases influencing their behaviour, and of remedies that could be undertaken to combat this. One model for thinking about these training sessions is in terms of individual responsibility: individuals are made aware of certain problematic facts (their propensity for bias), and are given certain tools that they can apply to their cognitions, or to their workplace procedures, in order to try to prevent their biases from having a role. This training devolves (forward-looking) responsibility to individuals for remedying bias.

However, this model may be flawed, both in terms of efficacy and theoretical warrant. First, if the emphasis is on individual de-biasing, this requires that

individuals accept and acknowledge that they themselves have implicit bias. Since cognitive biases include objectivity bias (the propensity to believe that one is more objective than one's counterparts (Pronin et al 2004; Rachlinski et al 2009)), there are significant obstacles to securing this acceptance. Moreover, we need not (and often cannot) pinpoint exactly whom is biased and on what occasions; it suffices to motivate institutional change to know that many individuals will, on many occasions, display certain kinds of biases. Instead, the emphasis should be on the fact that bias is pervasive, and as individuals within institutions we should seek procedures that make those institutions robust against the influence of implicit bias. Second, this model may seem to imply that the primary focus of addressing implicit biases is to de-bias individuals. If this is the aim, then devolving responsibility to individuals is a flawed strategy: since there exist few studies that show the long term efficacy of changes to implicit biases (Lai et al 2014, though cf. Madva ms.). The point is not that individuals should not try to de-bias (see Madva ms. for the claim that such efforts, are, in the grand scheme of things, rather small), nor that individual attitude change is not relevant at all – but that such efforts should not be independent of institutional change – which itself requires that individuals are motivated to institute and sustain those changes. Moreover, and fundamentally, this model risks obscuring the fact that individual de-biasing is only part of, and perhaps instrumental to, the central goal, which is to address patterns of discriminatory outcome. This may be secured by means other than individual de-biasing.

The overall point, then, is that in recognising the role that individuals and institutions must play, we see that institutional responsibility must reach further than simply providing implicit bias training, on the assumption that this devolves responsibility to individuals for dealing with discrimination. Rather, the responsibility

is with (individuals within) institutions to take sustained measures to address whatever mechanisms are producing discriminatory outcomes. Implicit bias training may be a part of those measures, but institutional change must extend well beyond this.

## References

- Amodio, D. & Devine, P., (2006), "Stereotyping and evaluation in implicit race bias: evidence for independent constructs and unique effects on behavior", *Journal of Personality and Social Psychology*, 91(4): 652
- Anderson, E. (2012). "Epistemic justice as a virtue of social institutions", *Social epistemology*, 26(2), 163-173.
- Blair, I. (2002) "The Malleability of Automatic Stereotypes and Prejudice", *Personality and Social Psychology Review*, 3: 242-261.
- Brownstein, M. (2015) "Attributionism and Moral Responsibility for Implicit Bias", *Review of Philosophy and Psychology* doi:10.1007/s13164-015-0287-7 pp.1-22
- Cameron, C. D., Payne, B. K., & Knobe, J. (2010). Do theories of implicit race bias change moral judgments? *Social Justice Research*, 23(4), 272-289. DOI: 10.1007/s11211-010-0118-z
- Cooley, E. Payne, B.K., Loersch, C., Lei, R. (2014) "Who owns implicit attitudes? Testing a metacognitive perspective", *Personality and Social Psychology Bulletin*, 41(1): 103-115
- Dovidio, J. F., Gaertner, S. E., Kawakami, K., & Hodson, G. (2002). "Why can't we just get along? Interpersonal biases and interracial distrust", *Cultural Diversity and Ethnic Minority Psychology*, 8(2), 88.



Eberhardt, J., P. Goff, V. Purdie, & P. Davies, 2004, "Seeing black: race, crime, and visual processing", *Journal of personality and social psychology*, 87(6): 876.

Faucher, L. (2016) "Revisionism and Moral Responsibility for Implicit Attitudes" in Brownstein, M., & Saul, J. (eds.) *Implicit Bias and Philosophy* volume 2 pp.115-145.

Fischer, J. & Ravizza, M. (2000) *Responsibility and Control: A Theory of Moral Responsibility* Cambridge University Press.

Follenfant, A. & F. Ric, 2010, "Behavioral Rebound following stereotype suppression", *European Journal of Social Psychology*, 40: 774–782.

Glasgow, J. (2016) "Alienation and Responsibility" in Brownstein, M., & Saul, J. (eds.) *Implicit Bias and Philosophy* volume 2 pp.37-61.

Green, A., D. Carney, D. Pallin, L. Ngo, K. Raymond, L. Lezzoni, & M. Banaji, 2007, "Implicit bias among physicians and its prediction of thrombolysis decisions for black and white patients", *Journal of General Internal Medicine*, 22: 1231–1238.

Greenwald, A., Poehlman, T.A., Uhlmann, E.L., Banaji, M.R., (2009) "Understanding and Using the Implicit Association Test: III. Meta-Analysis of Predictive Validity" *Journal of Personality and Social Psychology* 97: 17-41

Haslanger, S. (2015) "Distinguished Lecture: Social structure, narrative and explanation", *Canadian Journal of Philosophy*, 45(1), 1-15.

Holroyd, J. (2012) "Responsibility for Implicit Bias" *Journal of Social Philosophy*, Special Issue: Gender, Implicit Bias and Philosophical Methodology Crouch, M. & Schwartzman, L. (eds.) 43(3), pp. 274–306

Holroyd, J. (2014) "Implicit Bias, Awareness and Imperfect Cognition" in *Consciousness and Cognition*, Special Issue: Costs and Benefits of Imperfect Cognitions Bortolotti, L. & Sullivan-Bissett, E. (eds.) pp. 511-523

- Holroyd, J. (2016) "What do we Want from a Model of Implicit Cognition?"  
Proceedings of the Aristotelian Society 116(2) pp.153-179
- Holroyd, J & Kelly, D. (2016) "Implicit Bias, Character and Control" in Webber, J.  
and Masala, A. (eds.) From Personality to Virtue: Essays on the Philosophy of  
Character Oxford University Press, pp.106-134
- Holroyd, J. & Puddifoot, K. (forthcoming) "Implicit Bias and Prejudice" Routledge  
Handbook of Social Epistemology, edited by Miranda Fricker, Peter J. Graham, David  
Henderson, Nikolaj Pedersen, and Jeremy Wyatt.
- Jacobsen, A. (2016) "Reducing Racial Bias: Attitudinal and Institutional Change" in  
in Brownstein, M., & Saul, J. (eds.) Implicit Bias and Philosophy volume 2 pp.173-  
190
- Jost, J., Rudman, L., Blair, I., Carney, D., Dasgupta, N., Glaser, J. & Hardin, C.  
(2009) The Existence of Implicit Bias Is Beyond Reasonable Doubt: A Refutation of  
Ideological and Methodological Objections and Executive Summary of Ten Studies  
that No Manager Should Ignore Research in Organisational Behaviour 29: 39-69.
- Kang, J (2014) "Facing the Enemy: Segregation and Implicit Bias"  
[http://furmancenter.org/research/iri/essay/implicit-bias-and-segregation-facing-the-  
enemy](http://furmancenter.org/research/iri/essay/implicit-bias-and-segregation-facing-the-enemy) accessed 19-09-2016
- Lai, C. K., Marini, M., Lehr, S. A., Cerruti, C., Shin, J. L., Joy-Gaba, J. A., Ho, A. K.,  
Teachman, B. A., Wojcik, S. P., Koleva, S. P., Frazier, R. S., Heiphetz, L., Chen, E.,  
Turner, R. N., Haidt, J., Kesebir, S., Hawkins, C. B., Schaefer, H. S., Rubichi, S.,  
Sartori, G., Dial, C. M., Sriram, N., Banaji, M. R., & Nosek, B. A. (2014). "Reducing  
implicit racial preferences: I. A comparative investigation of 17 interventions",  
Journal of Experimental Psychology: General, 143, 1765-1785.

Legault, L., Gutsell, J.N., and Inzlicht M. (2011) “Ironic Effects of Anti-prejudice Messages: How Motivational Interventions can reduce (but also increase) prejudice”, *Psychological Science* 22(12) 1472-1477

Levy B. R., and Banaji M.R. (2002) “Implicit Ageism” in T.D. Nelson (ed.) *Ageism: Stereotyping and Prejudice Against Older Persons* Cambridge, MA.:MIT Press: 49-75.

Levy, N. 2012, “Consciousness, Implicit Attitudes, and Moral Responsibility”, *Noûs*, 48: 21–40.

Levy, N. 2014, “Neither fish nor fowl: Implicit attitudes as patchy endorsements”, *Noûs*. doi:10.1111/nous.12074

Levy, N., (forthcoming) “Implicit Bias and Moral Responsibility: Probing the Data” *Philosophy and Phenomenological Review*.

Madva, A. (2016a), “Virtue, Social Knowledge, and Implicit Bias”, in Brownstein & Saul (eds.) *Implicit Bias and Philosophy* volume 1.pp.191-215

Madva, A. (2016b) "A Plea for Anti-anti-individualism: How Over-simply Psychology Misleads Social Policy Ergo ms. available at: <http://alexmadva.com/sites/default/files/Madva%20-%20Plea%20for%20Anti-Anti-Individualism%202015.11.16.pdf> accessed 19-09-2016

Madva, A. (ms.), “Biased Against De-Biasing: On the Role of (Institutionally Sponsored) Self-Transformation in the Struggle Against Prejudice” <http://alexmadva.com/sites/default/files/Biased%20Against%20Debiasing%202015.4.pdf> accessed 19-09-2016

Malle, B. F., Guglielmo, S., & Monroe, A. E. (2014). “A theory of blame”, *Psychological Inquiry*, 25(2), 147-186

Mandlebaum, E. (2015) "Attitude Inference and Association: On the Propositional Structure of Implicit Bias" *Nous* 10.1111/nous.12089

Moskowitz, G.B., & Li, P. (2011). "Egalitarian goals trigger stereotype inhibition: A proactive form of stereotype control", *Journal of Experimental Social Psychology* Pp.172-179

Olson & Fazio 2006, "Reducing automatically activated racial prejudice through implicit evaluative conditioning", *Personality and Social Psychology Bulletin*, 32: 421–433.

Oswald, F.L., Mitchel, G. Blanton, H. Jaccard, J. Tetlock, PE., (2013) "Predicting ethnic and Racial Discrimination: A meta-analysis of IAT criterion studies" *Journal of Personality and Social Psychology* 105(2) 171-192.

Pickard, H. (2013) "Irrational Blame" *Analysis* 73(4): 613-626

Pronin, E. Gilovich, T., and Ross, L. (2004) "Objectivity in the Eye of the Beholder: Divergent Perceptions of Bias in Self Versus Others", *Psychological Review*, 111(3), Jul 2004, 781-799

Rachlinski, J.J., Johnson, S.; Wistrich, A.J., and Guthrie, C. (2009), "Does Unconscious Racial Bias Affect Trial Judges?" (2009). Cornell Law Faculty Publications. Paper 786. <http://scholarship.law.cornell.edu/facpub/786>

Saul, J. (2013a) "Implicit Bias, Stereotype Threat, and Women in Philosophy" in Jenkins, F. & Hutchinson, K. *Women in Philosophy: What needs to change?* Oxford University Press.

Saul, J. (2013b) "Implicit Bias and Scepticism" *Disputatio* 5: 37, 243-263.

Scaife, R. Holroyd, J., Stafford, T., Bunge, A., (ms.) "The Effects of Moral Interactions on Implicit Racial Bias" <https://osf.io/eubjp/>

- Sher, G. (2009) *Who Knew? Responsibility without Awareness*, Oxford University Press
- Sie, M., and Van Voorst Vader-Bours, N. (2016) "Stereotypes and Prejudices: Whose Responsibility? Indirect Personal Responsibility for Implicit Biases" in Brownstein, M., & Saul, J. (eds.) *Implicit Bias and Philosophy* volume 2 pp.90-114.
- Smith, A. (2005) Responsibility for Attitudes: Activity and Passivity in Mental Life," *Ethics* 115:2: 236-271.
- Smith, A. (2012) "Attributability, Answerability and Accountability: In Defence of a Unified Account" *Ethics* 122: 575-589
- Snow, N. (2005) "Habitual Virtuous Actions and Automaticity" *Ethical Theory and Moral Practice* 9 (5):545-561
- Stafford, T. (2014). "The perspectival shift: how experiments on unconscious processing don't justify the claims made for them", *Frontiers in Psychology*, 5, 1067. doi:10.3389/fpsyg.2014.01067
- Tetlock, P.F., Kristel, O., Elson, B., Green M., and Lerner, J. (2000) "The Psychology of the Unthinkable: taboo trade-offs, forbidden base rates and heretical counterfactuals" *Journal of Personality and Social Psychology* 78(5) 853-870
- Turner, M.A & Ross, S.L. (2007), "Housing Discrimination in Metropolitan America: Unequal Treatment of African Americans, Hispanics, Asian Americans, and Native Americans", in Goering, J. (ed.) *Fragile Rights Within Cities: Government, Housing and Fairness* Rowman & Littlefield Press
- Valian, V. (1999) *Why So Slow? The Advancement of Women* The MIT Press.

Vargas, M. (2017) "Implicit Bias, Responsibility and Moral Ecology"  
D. Shoemaker, (ed.), Oxford Studies in Agency and Responsibility Oxford University  
Press.

Vargas (2013) Building Better Beings: A Theory of Moral Responsibility. Oxford  
University Press.

Vargas (2009) "Revisionism about Free Will: A Statement and  
Defense" Philosophical Studies 144(1)pp. 45-62

Vargas (2005) "The Revisionist's Guide to Responsibility" Philosophical Studies 125  
(3), pp.399-429.

Washington, N. & Kelly, D. (2016) "Who's Responsible for This? Moral  
Responsibility, Externalism and Knowledge about Implicit Bias" in Brownstein, M.,  
and Saul, J. Implicit Bias and Philosophy Volume 2, Oxford University Press pp.11-  
36

Watson, G. (1996) "Two faces of Responsibility" Philosophical Topics 24(2): 227-  
248

Wood, M., Hales, J., Purdon, S., Sejersen, T. & Hayllar, O. (2009) A test for racial  
discrimination in recruitment practice in British cities: research report no 607.  
Department for Work and Pensions, London.

Wilson, M.C, and Scior, K. (2014) "Attitudes towards individuals with disabilities as  
measured by the implicit association test: a literature review". Research in  
Developmental Disabilities; 35(2): 294-321.

WISELI (2012) Reviewing Applicants: Research on Bias and Assumptions 3rd  
Edition [http://wiseli.engr.wisc.edu/docs/BiasBrochure\\_3rdEd.pdf](http://wiseli.engr.wisc.edu/docs/BiasBrochure_3rdEd.pdf) accessed  
19/09/2016

Zheng, R., (2016), "Attributability, Accountability and Implicit Attitudes", in  
Brownstein & Saul (eds.) *Implicit Bias and Philosophy* vol.2 pp.62-89