



This is a repository copy of *What Do We Want from a Model of Implicit Cognition?*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/98670/>

Version: Accepted Version

Article:

Holroyd, J.D. (2016) What Do We Want from a Model of Implicit Cognition? Proceedings of the Aristotelian Society, 116 (2). pp. 153-179. ISSN 0066-7374

<https://doi.org/10.1093/arisoc/aow005>

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 including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

What Do We Want from a Model of Implicit Cognition?

Journal:	<i>Proceedings of the Aristotelian Society</i>
Manuscript ID	Draft
Manuscript Type:	Invited Paper
Date Submitted by the Author:	n/a
Complete List of Authors:	Holroyd, Jules; University of Sheffield, Department of Philosophy
Keywords:	implicit bias, implicit cognition, belief, alief, automaticity, racism

SCHOLARONE™
Manuscripts

Manuscript Only

What do we want from a model of implicit cognition?

Jules Holroyd

Abstract: In this paper, I set out some desiderata for a model of implicit cognition. I present test cases and suggest that, when considered in light of them, some recent models of implicit cognition fail to satisfy these desiderata. The test cases also bring to light an important class of cases that have been almost completely ignored in philosophical discussions of implicit cognition and implicit bias. These cases have important work to do in helping us understand both the role of implicit cognition in action, and our attempts to combat implicit biases.

In this paper, I set out some desiderata for a model of implicit cognition. I present test cases, and suggest that, when considered in light of them, some recent models of implicit cognition and in particular, of implicit bias - from Levy, Schwitzgebel, Mandelbaum, Gendler and Machery - fail to meet these desiderata. I propose a 'minimal model' that may fare better. I also bring to light an important class of cases that have been almost completely ignored in philosophical discussion of implicit cognition and implicit bias. These cases have important work to do in helping us understand both the role of implicit cognition in action, our responses to implicit bias and attempts to combat it.

Models of implicit cognition have emerged from different sources with different priorities and conceptual frameworks. Psychologists have developed such models as a way of understanding what processes or mechanisms best explain the experimental findings about implicit associations and their impact on behaviour. Meanwhile, recent philosophical accounts of implicit cognition have focused on implicit bias in particular. These accounts have been formulated with various aims in mind, including: incorporating implicit bias into a complete picture of our psychological states and their role in action; articulating the role of these implicit cognitions in a normative conception of agency or responsibility; identifying the epistemic implications of this aspect of our agency. To this end, it is no surprise that philosophical conceptions of implicit bias vary significantly. But how might we go about evaluating whether a model of implicit bias is successful? And how do the accounts developed by philosophers, to date, fare? My task in this paper is to start to address these questions.

In section I, I set out some desiderata for a model of implicit cognition. The desiderata I set out are not intended to be exhaustive, but provide a starting framework for evaluating models of this aspect of our cognition. In section II, I introduce some test cases which can be used to evaluate whether the accounts of implicit bias satisfy our desiderata. These cases also bring to the forefront some cases of implicitly biased action that have been almost wholly absent from the philosophical discussion of implicit bias. In section III, I then evaluate five views which provide models for understanding what some of our implicit cognitions - implicit biases - are. Levy claims that implicit biases are patchy endorsements; Schwitzgebel claims they are in-between beliefs; Mandelbaum suggests they are unconscious beliefs; Gendler models them as *aliefs*; Machery argues that implicit attitudes properly understood are not implicit at all. I evaluate how well each view is able to satisfy the desiderata from section I, and find them wanting. They cannot accommodate our test cases from section II, nor capture the distinctive explanatory and moral significance of implicit cognition that we observe in these cases. I close with a proposal for an account that might more satisfactorily meet these desiderata.

I

Desiderata for a model of implicit cognition. The sort of phenomena that models of implicit bias try to capture are by now familiar: the processes or states that have a distorting influence on behaviour and judgement, and are detected in experimental conditions with implicit measures, such as the

1
2 implicit association test. The most troubling of these concern distortions to do with social identity:
3 gender, race, age or sexuality biases, for example. Such implicit cognitions are typically understood
4 as fast and automatic processes, which may operate without the agent being aware of their
5 activation or influence on behaviour. For example, implicit racial biases have been found to have a
6 role in influencing the evaluation of CVs - participants rated identical CVs less positively when the
7 name at the top indicated that it was the CV of an individual racialised as black (Dovidio &
8 Gaertner 2000). Implicit biases may manifest in 'micro-behaviours' - unintentional non-verbal
9 behaviours that indicate tension or discomfort. White people are found to display these in interracial
10 interactions with black interlocutors, affecting the quality of the interpersonal interactions (Dovidio,
11 Kawakami & Gaertner 2002; Dovidio, Gaertner, Kawakami & Hodson 2002). And, disturbingly,
12 implicit associations between black people and weapons are found to affect performance on shooter
13 bias tasks, with individuals more likely to make the error of shooting an unarmed individual when
14 that individual is black (Glaser & Knowles 2007).

15
16 Since these judgements and behaviours are poorly described by more familiar psychological
17 states - such as the agent's explicit beliefs, intentions and values - we need an account of what it is
18 that produces such discriminatory behaviour. Accounts of implicit bias aim to capture the parts of
19 our psychology implicated in these behaviours. Such an account aims to articulate what this part of
20 our psychology consists in; what it does; and how it is integrated with other cognitions. The inquiry
21 is complicated by the fact that there is no consensus regarding what it is that makes a part of our
22 mental workings 'implicit' rather than explicit. It is also complicated by the fact that a broad range
23 of phenomena may fall under the rubric of implicit cognition, and it may be that there are different
24 understandings of 'implicit' involved in these phenomena. Our present focus is on models of
25 implicit bias. In evaluating accounts of implicit bias, we should at this stage leave open both the
26 question of what is distinctively 'implicit' about such cognitions, and whether an adequate account
27 of implicit bias is continuous with other aspects of our implicit cognitions. With this in mind, we
28 begin by asking what we want such an account to do.

29
30
31
32 *(i) Empirical validity.* Our theory should be sensitive to the findings of empirical psychology. A
33 wealth of empirical data can be drawn on to construct our theory of implicit cognition. For example
34 one dimension of the debate has focused on whether implicit attitudes have propositional or
35 associative structure (Mandelbaum 2015, Levy 2015). The theory best placed will be (inter alia) the
36 one best supported by the empirical evidence about whether implicit biases seem to operate in a
37 way that only states with propositional rather than associative structure do - for example, whether
38 they appear to be implicated in inferential patterns of reasoning or not. More generally, we should
39 expect that the preferred model of implicit cognition finds support in the empirical literature. There
40 are two dimensions of empirical validity in particular that are worth highlighting, as follows.

41
42
43 *(ii) Bias in context.* We know that implicit biases are, in part, caused by the associations we find in
44 our cultural and social environment - such that race biases are contingent upon a specific cultural
45 context that encodes racist stereotypes (Devine 1989). We also know that implicit biases can be
46 reigned in by institutional structures, and biases unlearned (or new, non-stigmatising biases learned)
47 in the appropriate social contexts.¹ Accordingly, our model of implicit cognition should cohere with
48 these findings and should elucidate our understanding of how implicit cognitions are shaped by
49 these social and institutional structures.

50
51
52 *(iii) Framework for understanding how to change implicit cognitions.* Since our focus is on implicit
53 biases that concern the stigmatisation of social identities, one important aspect of a theory of
54 implicit cognition is that it should help us to understand how we might change or combat implicit
55

56
57
58 ¹ See Jost et al 2009 for an overview of institutional contexts in which biases might operate; and Lai et al 2014, for changes that
59 may prevent biases from operating. See also 'Beyond Bias and Barriers' by the NAC for specific proposals for institutional
60 change (in the context of STEM research).

1 biases. The theory should deliver hypotheses that generate testable predictions about what might
 2 change implicit biases, or how their expression might be mitigated. This is not to say that the model
 3 must deliver prescriptions for changing implicit biases; but rather that a useful model will generate
 4 hypotheses that we can test. It is consistent with this that when tested, no interventions are found to
 5 be successful. This at least helps us to see where we should not allocate our efforts. The point is that
 6 if a theory can generate predictions, the confirmation of those predictions provides additional
 7 empirical validity. Moreover, it is desirable that a theory of implicit bias aid us with the
 8 transformative work needed to remedy problematic implicit biases.
 9
 10

11 However, the empirical findings are often consistent with more than one interpretation of the data,
 12 and under-determine which model we might endorse. So we need further desiderata to guide us.
 13
 14

15 *(iv) Coherence with our understanding of cognition in general.* We should want our model of
 16 implicit bias to integrate with our best understanding of cognition in general. This is not to say that
 17 our theories of implicit cognition should be constrained to utilise only those notions or posits of
 18 extant theories of mind. It may be (we shall indeed see examples of this shortly) that new sorts of
 19 state, or revisionary understandings of existing states, have to be posited in order to make good
 20 sense of the implicit workings of our psychologies. But we should be in a position to explain how
 21 those workings can be integrated into a fuller picture of our minds. One hope is that we get an
 22 account of implicit bias that is able to explain how implicit biases relate to, and interact with,
 23 explicit cognitive processes or explicit mental states. This requires also that our model elucidates
 24 what it is for this aspect of cognition to be *implicit* and how it is distinguished from explicit aspects
 25 of cognition. This is no small task, of course - and as we shall see, there are problems with extant
 26 understandings of this.
 27
 28

29 *(v) Framework for understanding the role of implicit cognition in moral agency.* Since many
 30 philosophers engaged in this debate are not only interested in what is in our minds, and how that
 31 influences behaviour, but also for what this means for beings who are moral agents, we should seek
 32 a theory of implicit cognition that coheres with a framework for understanding how this aspect of
 33 our cognitive workings is implicated in our moral agency. This is not to say that the theory should
 34 be constrained by folk psychological notions of what moral agency consists in. Moreover, we
 35 should remain open to our judgements about moral agency and moral responsibility needing
 36 considerable refinement or revision in order to be suitably responsive to empirically informed
 37 understandings of our agency. It may be that our best account - along with assumptions about the
 38 nature of moral agency, responsibility and evaluation - delivers the judgement that implicit biases
 39 are not the sort of thing for which we can be morally evaluated. But our understanding of moral
 40 agency and moral responsibility provides some useful points of reference for evaluating models of
 41 implicit bias.² Suppose our judgements about agents' responsibility clearly distinguishes morally X
 42 and Y. But our model of implicit bias fails to distinguish between the moral status of X and Y. This
 43 gives us some reason both to consider revising our judgements about X and Y, and to check whether
 44 our model of implicit bias rests on defensible assumptions, proper interpretation of empirical
 45 findings, and so on.³
 46
 47
 48

49 There may be other desirable features of a model of implicit cognition. I do not claim these to be
 50 exhaustive. However, I take these to be useful desiderata when evaluating such models. I next
 51
 52

53
 54 2 In articulating the desiderata for an analysis of discrimination, Lippert-Rasmussen suggests that an analysis should be 'morally
 55 enlightening', in helping us to tease out moral conclusions (if not entailing them) (2014, 48). This is what I have in mind here.

56 3 There are some respects in which the project of articulating desiderata resonates with Cudd's (2005) articulation of desiderata
 57 with which to evaluate a theory of oppression. However, note that my desiderata differ in some important respects: unlike
 58 Cudd's criteria of axiological pragmatism (p.32), I do not require that the theory tell us how to change problematic biases (but
 59 merely generate predictions for this). Likewise, the desiderata regarding a moral framework does not make any claim about the
 60 primacy of the value of moral agency (as Cudd does, p.36).

1 delineate a set of test cases. These cases help us to see the importance of these desiderata; second,
 2 they bring to light a kind of case that has been neglected in the philosophical literature; and thirdly,
 3 we can use these cases to test whether the models of implicit bias meet these desiderata, in section
 4 III.
 5
 6

7 II

8 Test cases. A notable feature of the philosophical literature on implicit bias is that it has focused
 9 overwhelmingly on what I call ‘conflict’ cases of implicit biases. By this I mean cases in which an
 10 agent’s explicit values, commitments or endorsements are in conflict with a posited implicit bias
 11 which influences her actions in ways contrary to those explicit values, commitments and
 12 endorsements.
 13

14 For example Saul proceeds with her attention on implicit biases in ‘those who explicitly and
 15 sincerely avow egalitarian views’ (2013, p.41). Glasgow’s discussion of responsibility and
 16 alienation takes as its main focus cases of implicit biases which are in conflict with the agent’s
 17 endorsed values (2016) - hence the alienation. Washington and Kelly’s discussion focuses on cases
 18 in which there is ‘dissociation’ between implicit and explicit attitudes (2016). And Levy’s analysis
 19 of implicit attitudes centres on cases in which implicit attitudes come apart from the agent’s
 20 endorsed explicit attitudes, and influence behaviour in unwanted ways (2014, forthcoming).
 21 Conflict cases have been the overwhelming focus of the philosophical literature.
 22

23 There are good reasons for focusing on these cases. Firstly, they are startling to many of us,
 24 since we have fair-minded explicit values, and commitments to non-discrimination. So it is
 25 discomfoting to find that we have implicit biases that might implicate us in discrimination.
 26 Secondly, pragmatically speaking, these cases are an important class to focus on - for those with
 27 explicitly anti-discriminatory values and commitments, such as ourselves, are likely to be well
 28 motivated to endorse and implement efforts to overcome the influence of implicit biases. And
 29 finally, conflict cases raise a host of interesting and important philosophical considerations: what is
 30 the moral status of such implicit biases since they conflict with professed evaluative endorsements?
 31 Is the agent responsible for them? And so on.
 32

33 However, this focus misses out an important class of cases. The cases I have in mind are
 34 those in which agent’s implicit biases are aligned with her explicit attitudes. Whilst these cases have
 35 not received much philosophical attention, they are by no means mere conceptual possibilities, but
 36 empirical realities. Indeed, we know that agents with greater explicit prejudice record stronger
 37 implicit biases on implicit measures: Devine et al (2002) found that implicit race biases were
 38 stronger in individuals who harboured more explicit racial prejudice (see also Glaser & Knowles
 39 2007). Accommodating non-conflict cases is important if a theory is to meet our desiderata.
 40 Consider the following three cases of individuals and their actions:
 41
 42

43 Racist. Individual A has explicitly prejudiced beliefs and negative evaluations about black
 44 people. Further (and in accordance with what would be predicted on the basis of Devine et
 45 al’s (2002) studies) this individual also harbours implicit negative associations. When on a
 46 hiring committee, A’s evaluations are guided by his explicit prejudices (he does not want a
 47 black colleague, he believes that race is relevant to the suitability of a candidate). His
 48 behaviour is also influenced by his implicit biases, and he unintentionally displays certain
 49 ‘micro-behaviours’ in interracial interactions: he manifests discomfort, reacts with more
 50 irritability, sits marginally further away from the black interviewees (Dovidio, Kawakami,
 51 Gaertner & Hodson 2002, McConnell & Leibold 2001). He is not aware of these aspects of
 52 his behaviour, though he would not disavow them if he were. Unsurprisingly, A’s makes an
 53 intentionally discriminatory decision: between the two equally well qualified candidates,
 54 the white rather than black candidate is hired.
 55
 56

57 Protocol adhering racist. Individual B has explicitly prejudiced beliefs and negative
 58 evaluations about black people. Further (and in accordance with what would be predicted on
 59
 60

1 the basis of Devine's studies) this individual also harbours negative implicit associations. B
 2 also is strongly motivated in his professional role, and is very concerned to adhere to
 3 protocol, including fair hiring procedures. And indeed, if the best person to do the job is
 4 black, he thinks, they should be hired (still, his racism manifests in his preferences against
 5 going for after work drinks or invite anyone 'like that' round for dinner). B forms the
 6 intention not to discriminate when given responsibilities on a hiring committee. B's
 7 evaluations of the applicants are guided by his explicit intention not to discriminate.
 8 However, his behaviour is also influenced by his implicit biases: like A, he displays
 9 unintentional micro-behaviours: he manifests more discomfort, reacts with more irritability,
 10 and sits marginally further away from the black interviewees. He hasn't reflected much and
 11 is not aware of these aspects of his behaviour. Whilst these unintentional aspects of
 12 behaviour aligned with his explicit racist attitudes, they do not accord with his intention to
 13 adhere to protocol. These micro-behaviours affect the quality of the interaction, and when
 14 deciding between the two best qualified candidates, on the basis of the interview
 15 performance inflected by B's implicit biases, B judges there to be a clear margin between
 16 them - and the equally well qualified black candidate is not hired.

17
 18
 19
 20
 21 *Biased egalitarian:* This by now familiar individual, C, has explicitly anti-racist beliefs and
 22 egalitarian commitments. However, this individual also harbours implicit negative
 23 associations. When on a hiring committee, he intends not to discriminate (in accordance
 24 with his explicit anti-racist commitments). However, his implicit race associations mean
 25 that, like A and B, his behaviour is influenced by his implicit biases, and he manifests
 26 various micro-behaviours: he displays more discomfort, reacts with more irritability, and sits
 27 marginally further away from the black interviewees. These micro-behaviours affect the
 28 quality of the interaction, and when deciding between the two best qualified candidates, the
 29 black candidate is not hired.
 30

31
 32 Considering these three cases alongside each other helps us to see the role of our desiderata. First,
 33 an empirically valid account of implicit bias should be able to accommodate both conflict (C) and
 34 non-conflict cases (A and B), if they are to meet the desiderata of empirical validity: in each of
 35 these cases, implicit bias is implicated in discriminatory behaviour. Secondly, these cases bring out
 36 the differential first pass moral evaluations we might make of each individual. Each individual
 37 discriminates, but the causal and moral story behind this differs: A has morally repugnant attitudes
 38 and values, and intentionally discriminates. B too has repugnant explicit attitudes, but the fact that
 39 his discrimination is not intentional makes an important difference in the judgement of his
 40 behaviour.⁴ C marks the paradigm 'conflict case' with which we are familiar from the philosophical
 41 literature: it is the conflict between the commitments and the discriminatory actions due to implicit
 42 bias that has garnered so much philosophical intrigue and attention. C too unintentionally
 43 discriminates. Some might not see a great deal of moral difference between agents in B and C,
 44 insofar as both unintentionally discriminate. But there are reasons to suppose that the agent in B is
 45 guilty of moral failures that the agent in C is not. Aside from the repugnant moral views B holds, he
 46 is also guilty of further moral and epistemic failures. Like A, even in the absence of knowledge
 47 about implicit bias, it is reasonable to expect that racist explicit beliefs will inflect automatic aspects
 48 of behaviour. B is guilty of culpable ignorance, or negligence in supposing that his behaviours
 49 would not be inflected by these morally problematic attitudes. On the other hand, an important
 50 difference between A and B is that whereas A would endorse these micro-behaviours were he aware
 51 of them, it is less clear that B would do so - despite their convergence with his explicit racism, they
 52
 53
 54

55
 56
 57
 58
 59
 60
 4 I am assuming that implicit biases have a role in B's action. Why think this, rather than that a different mechanism is at work than in C, where there is no explicit prejudice? Since we find implicit bias prior to, and in the absence of, explicit biases, and since we know that various automatic cognitions run alongside our deliberative and explicit cognitions, it would not be parsimonious to posit a different mechanism in the case of B, and to suppose that implicit bias does not have any role in the discriminatory outcome here. I thank Zoe Drayson for pressing me on this point.

thwart other explicit intentions he holds. So, we can see the importance of a theory of implicit cognition delivering a framework that can make sense of these differential judgements.

These test cases also drive home the importance of a theory that can provide hypotheses about differential interventions regarding implicit biases. Restricting our attention to interventions that would tackle the racist discrimination manifested in each context, we can see that each case may require different interventions. For individual A, the racist, we have to address explicit as well as implicit racism; for the protocol adhering racist, B, strategies are needed that help to stop the bias taking effect (given the likelihood that the bias might persist due to its accordance with explicit prejudices); for the biased egalitarian, strategies would be welcome that insulate from the bias, and presumably the individual would also welcome strategies that change the implicit association itself.

III

Extant theories. In this section, we use these test cases to evaluate whether extant models of implicit bias can satisfactorily meet the five desiderata.

(i) Patchy-endorsements. Let us first consider Levy's account of implicit attitudes. Locating his view in the context of the dispute over whether implicit attitudes are propositional or associative in structure, Levy argues that they are best understood as 'patchy-endorsements'. That is, they are mental states which may have some propositional structure, but are not responsive to reasons in the way that explicit beliefs typically are.⁵ What follows from this, crucially, is that implicit attitudes do not stand in inferential relations with other attitudes that are attributable to the agent. They do not stand in these relations because of their 'patchy' structure, which means they are not inferentially integrated with a wide range of other attitudes that are attributable to the agent. There is evidence that at least sometimes, implicit attitudes are acquired notwithstanding their conflict with the agent's endorsed attitudes, and are not extinguished when this is recognised (Levy, forthcoming, p.19). Accordingly, there is reason to believe that implicit attitudes are not integrated in the right way with agent's other attitudes. This is important, Levy claims, because attitudes that are attributable to the agent make up her deliberative standpoint, and those attitudes central to this standpoint reflect her evaluative position. Accordingly, for Levy, agents should not be held accountable for action that is influenced by implicit attitudes insofar as it is true that 'had the agent's explicit attitudes controlled behaviour, the actions would have had a different moral character' (forthcoming, p.25).⁶

There are two important claims from Levy to keep in mind: that because implicit attitudes are patchy, they will not be well-integrated with the agent's endorsed attitudes; and secondly, that insofar as they are not integrated in this way, the moral character of action influenced by implicit attitudes will diverge from the moral character of actions had her explicit attitudes guided them. Recalling our desiderata, we can say that at first glance, Levy gives us a theory supported by empirical evidence about the acquisition and extinction of implicit attitudes, and the extent to which they are responsive to reasons. This view locates implicit attitudes in terms of extant posits in philosophy of mind, whilst holding such attitudes importantly distinct from them. This theory also provides a framework for making sense of the moral evaluations we might make of implicit attitudes when they influence action: a framework which exculpates agents when the counterfactual is true. We also get some understanding about how to change implicit attitudes: we cannot suppose that propositional reasoning will be successful, since these attitudes are patchy in their responsiveness to reasons.

What does Levy's view say about our set of test cases above? The crucial two cases for our

-
- 5 Levy notes that implicit attitudes might meet Fischer & Ravizza's (2000) conditions for reasons-responsiveness - namely, that a mechanism is moderately responsive to reasons (responsive in a somewhat patterned way across some possible worlds). He takes this as a reason to reject that standard as sufficient. I am not convinced by that, but grant it insofar as it motivates us to look more closely at the issue of attributability.
- 6 Earlier, Levy's statement of the counter-factual differs slightly: 'had the agent's explicit attitudes alone controlled their behaviour, the action would have lacked its actual moral character' (7). I think this statement of the counterfactual generates slightly different judgements - but it is the version stated in the text that Levy relies upon in his attention to non-conflict cases, so I follow him in this.

1 purposes are those in which the agents implicit and explicit attitudes are aligned (individual A and
 2 B). Levy himself does not say much about such cases, since his focus is mainly on conflict cases,
 3 where the implicit and explicit attitudes diverge. But he does remark that, in cases in which implicit
 4 and explicit attitudes are aligned (where the implicit attitude is ‘annexed’ to the agent and ‘well
 5 integrated’ with her endorsed attitudes) the counterfactual claim we saw above will be false. That is
 6 to say it will be false that ‘had the agent’s explicit attitudes controlled behaviour the action would
 7 have had a different moral character’ (p.25).
 8

9 We can already start to see that something is going wrong here. Consider, from our set of
 10 test cases above, the case of the protocol-adhering racist (individual B). Here we have an individual
 11 who (ex hypothesi) has aligned implicit and explicit attitudes: his implicit and explicit attitudes are
 12 both racist. His implicit attitudes influence his behaviour. On Levy’s view, since the implicit
 13 attitudes are aligned with the agent’s explicit attitudes, it is false that had B’s explicit attitudes
 14 controlled behaviour, the action would have had a different moral character. But this is wrong: had
 15 B’s explicit attitudes controlled the behaviour, it *would* have had a different moral character. As
 16 described, the action is one of unintentional discrimination due the implicit biases influencing
 17 behaviour. Had his explicit racist attitudes controlled behaviour, it would have been an intentionally
 18 discriminatory action. And had his explicit concern for protocol controlled his behaviour, it would
 19 have been a non-discriminatory action. Both of these are importantly different in moral character
 20 from an action that is unintentionally discriminatory. In the case of individual B, the relevant
 21 counterfactual claim at issue is true: the action *would* have a different moral character if guided by
 22 explicit attitudes. But we might nonetheless see such attitudes as attributable to the agent, and ones
 23 he is properly accountable for. So the counterfactual specified by Levy is not a helpful (part of a)
 24 moral framework for evaluating morally agents, their implicit and explicit attitudes, and actions.
 25

26 The case of individual B also helps us to see that it is a mistake to emphasise the lack of
 27 integration of implicit biases as a definitive feature of their characterisation. By identifying implicit
 28 biases as patchy endorsements - as states that, by virtue of their characterisation as ‘patchy’, are not
 29 well integrated with the agent’s explicit attitudes - Levy marginalises non-conflict cases of implicit
 30 attitude. But this is a mistake, insofar as implicit biases still have an explanatory role in our
 31 understanding of the full scope of discriminatory behaviour in such cases. In the case of individual
 32 A (the racist), there will be some subtle automatic or autonomic responses that are not under the
 33 control of the agent’s explicit attitudes (his fear responses, increased sweating and blink rates or
 34 dilated pupils, for example (Dovidio et al 1997)). In the case of individual B (the protocol-adhering
 35 racist) there will likewise be these aspects of his behaviour which, whilst supported by his explicit
 36 racist attitudes, are divergent from his explicit intention to adhere to protocol and avoid
 37 discrimination in hiring. Identifying these aspects of behaviour is crucial both for adequate moral
 38 evaluation, and for being able to target interventions to tackle discrimination appropriately.
 39

40 In sum, it is unsatisfactory to characterise implicit attitudes as patchy endorsements which
 41 are necessarily poorly integrated with and divergent from the agent’s explicit attitudes.⁷ This
 42 characterisation omits, by conceptual fiat, cases of aligned implicit and explicit attitudes. And this
 43 omission is problematic insofar as it fails to provide adequate resources for moral evaluation, nor
 44 for complete prescriptions about interventions to prevent discriminatory actions.
 45

46
 47
 48 *(ii) In-between beliefs.* Schwitzgebel (2010) provides us with a different model for thinking about
 49 implicit biases: as cases of in-between belief. The idea behind this picture is that, on a broad track
 50 dispositional model of belief,⁸ there will be cases in which it is most appropriate to say that we
 51 ‘kind of’ believe some proposition; our cognitive position is pretty much in-between. As
 52 Schwitzgebel illustrates, we suppose an individual has some of the dispositions associated with
 53 holding some belief, P: believing that family life is more important than work. She might avow P,
 54

55
 56
 57 7 Indeed, we could think that biases can be aligned whilst not being well-integrated, in the sense of not standing in the appropriate
 inferential or reasons-responsive relations. But this distinction is elided in Levy’s claims.

58 8 Schwitzgebel notes that the idea of in-between belief can also be rendered cogent on other models of belief, functionalist,
 59 representationalist, and so on.
 60

1
2 act as if P in context C1 (for example, prioritising school pick up). But the agent also has other
3 dispositions that are indicative of holding the belief not-P (that it is not the case that family life is
4 more important than work). For instance, the agent might act in context C2 as if not-P were true
5 (working late, drifting into work thoughts whilst on family holiday, and so on). The idea is that in
6 such cases, it is inappropriate to suppose that the agent flat out believes P or not-P; nor to ascribe
7 both beliefs; nor to posit that she switches back and forth between P and not-P. Rather, she *kind of*,
8 or in-between believes, both.
9

10 Implicit biases can be understood on this model, Schwitzgebel suggests. (It is worth noting
11 that he offers implicit biases as a case in support of his claims about in-between belief, rather than
12 as claims about a theory of implicit biases. Nonetheless, we can still ask whether this model really
13 is a useful way of thinking about implicit biases.) An individual, such as C from our test cases, has
14 some dispositions that we stereotypically associate with egalitarian commitments to fair treatment
15 and anti-discrimination, and beliefs in racial equality. She may avow such beliefs, ensure her
16 workplace has policies to address discrimination, attend anti-racism activist events. But an
17 implicitly biased individual may also have and manifest dispositions associated with the racist
18 belief that differently racialised individuals are not equal; she may react with greater hostility in
19 interracial interactions or express surprise when discovering the academic excellence of one of her
20 black students. These are dispositions stereotypically associated with racist beliefs or attitudes. In
21 such cases we say that the agent in-between believes in racial equality. Some of her dispositions are
22 consistent with this belief; others not - so we should not ascribe the belief wholesale, nor its
23 contrary to her; rather, she *kind of*, or in-between believes, in racial equality.
24

25 This model of in-between belief seems helpful when we consider conflict cases of implicit
26 biases - cases where the biases are in tension with the agent's explicit beliefs and values, or thwart
27 her pursuit of them. And, it does so in a way that seems to satisfy some of our desiderata - we find
28 implicit biases unpacked in terms that cohere with extant posits in our theories of cognition (beliefs,
29 understood in dispositional terms). This provides us with a framework in which we can start to
30 think about moral evaluation - to the extent that we have views about responsibility for belief, we
31 can consider importing those to apply to these cases of in-between belief (to what extent are the
32 agent's dispositions attributable to her, reflective of her evaluative stance, and so on). We see this
33 aspect of the view in Schwitzgebel's remarks about responsibility for implicit bias (he argues
34 individuals are responsible insofar as implicit biases are attributable to the agent, and reflective of
35 her personality and character).⁹
36

37 However, this analysis is not so helpful when it comes to understanding the role of implicit
38 bias in non-conflict cases - cases where some of the agent's behavioural dispositions (those related
39 to implicit biases) line up with the agent's explicit attitudes and beliefs (in particular in case A
40 above). In such cases, there is clearly no issue of 'in-between belief', since the kind of behavioural
41 dispositions (due to implicit cognition) that we find line up pretty neatly with the agent's explicit
42 beliefs or attitudes. Whilst not strictly speaking incompatible with recognising the explanatory
43 significance of implicit biases, adopting in-between belief as a model of implicit bias obscures the
44 role implicit bias can play in cases where the implicit and explicit attitudes are aligned. This is
45 especially so since, according to Schwitzgebel, we ascribe beliefs based on the observed
46 dispositional profile of the agent. And it is the dispositions 'features of which are broadly
47 recognisable to normal attitude ascribers' (2010, p.12) that determine whether a belief or attitude is
48 ascribed. Consider how this goes in conflict cases: we find that the agent behaves 'as if' they had
49 racist beliefs despite their explicit avowals, and the behavioural outcome itself leads us to ascribe
50 the in-between belief. But in non-conflict cases, such as that of individual A, we find the
51 behavioural outcome is explainable by the agent's intentional decision to discriminate. The problem
52 with this move is that the various dispositions associated with harbouring implicit race biases
53 include also subtle differences in interpersonal interactions, autonomic and automatic responses, as
54 described above. Yet insofar these behavioural dispositions are not part of the repertoire of 'normal
55
56
57

58
59
60

⁹ <http://www.newappsblog.com/2015/03/on-being-blameworthy-for-unwelcome-thoughts-reactions-and-biases.html>

1 attitude ascribers', they may be overlooked when there are intentional discriminatory decisions that
2 explain the agent's behaviour. Accordingly, it is not obvious that the dispositions that manifest in,
3 for example, unintentional micro-behaviours will be given due role in explaining behaviour in cases
4 of aligned implicit and explicit attitudes. Whilst some of the fine-grained behavioural dispositions
5 (the autonomic responses) are well captured in conflict cases, as 'in-between belief'; they are not
6 aptly described as such in these non-conflict cases. This obscures the role of the dispositional
7 profile associated with the implicit aspects of cognition.
8

9 I don't want to overstate the case against the in-between belief model of implicit attitudes.
10 The claim isn't that there are no resources to make sense of the sort of non-conflict cases I have
11 described. Rather the worry is that the role of implicit bias in such cases is easily obfuscated or
12 obscured on this account, and we should be alert to these difficulties in order to deliver a more
13 satisfactory model.
14

15
16 *(iii) Unconscious beliefs.* The last two views posited revisionary states to explain implicit biases:
17 patchy-endorsements and in-between beliefs. On Mandelbaum's view, such revisionary states are
18 not needed: rather implicit biases are best understood as 'honest-to-god propositionally structured
19 mental representations that we bear the belief relation to' (2015, p.7). Unlike the sort of beliefs we
20 are familiar with, though, implicit biases are unconscious beliefs. This means that (not necessarily,
21 but at least typically) they do not figure in our conscious cognitions, but can function inferentially,
22 and in reasons-responsive ways, beyond the reach of our reflective awareness. On this view, whilst
23 a person such as individual C may have an explicit belief, for example, that *race is irrelevant to the*
24 *quality and suitability of an applicant*, she may nonetheless harbour an unconscious racist belief,
25 with the propositional content *black people are less likely to be suitable candidates*.
26

27 This view is primarily motivated by a swathe of research findings that are difficult to make
28 sense of on the assumption that implicit cognition is purely associative (the sort of findings that also
29 inform Levy's commitment to the claim that implicit attitudes may have propositional content,
30 albeit (for Levy) content that is inferentially sensitive in a patchy sort of way). For example,
31 Mandelbaum canvases studies in which implicit cognition seems to respect the dictum 'the enemy
32 of my enemy is my friend' - difficult to make sense of on a purely associative model, where one
33 might expect negative associations to accrue to the enemy of my enemy (two relations of dislike
34 being at issue). Instead positive associations with the target object 'the enemy of the enemy' are
35 recorded, which suggests sensitivity to content (Gawronski et al 2005).
36

37 Accordingly, this view stakes a claim to meeting the first desiderata of empirical validity.
38 Moreover, as emphasised by Mandelbaum, this model crucially helps us to predict the sorts of
39 strategies worth evaluating for efficacy in mitigating implicit bias. Those which focus only on the
40 counter-conditioning of associative states will fail to exploit the extent to which implicit biases may
41 be sensitive to inferential interventions. Likewise, it is also pretty clear that the posits of this view
42 cohere with those of our extant theories of cognition, so the unconscious belief model of implicit
43 bias fares well with respect to this desiderata, at least at first blush. Finally, insofar as we have
44 available to us frameworks for understanding the moral evaluation of beliefs and actions informed
45 by them (see for example Hieronymi, 2008), this view appears to stand well with respect to the our
46 desiderata regarding the framework for moral evaluation of implicit biases.
47

48 How does this view fare in dealing with our test cases, however? We might first worry
49 whether the unconscious belief account will capture all aspects of the phenomenon described in our
50 test cases, which include the activation of affective and autonomic responses as well as responses
51 that may more aptly be described as belief. At best this account has to see these as the causal upshot
52 of the unconscious beliefs. Even with that caveat, the difficult cases for Mandelbaum's view are
53 those in which there is alignment between the implicit and explicit biases. How might the
54 unconscious belief model handle these cases? One option would be to say that there are two beliefs
55 here: one explicit (the racist belief), the other unconscious (with the same content). In case B (the
56 protocol-adhering racist) the unconscious belief is causally efficacious even whilst the explicit one
57 is not, in case (A) both beliefs are causally efficacious. This is problematic, insofar as it commit us
58
59
60

1 to the idea that beliefs are individuated, and distinguished, not by their content (or indeed functional
 2 role) but according to whether they are conscious or not (hence we have two beliefs here, rather
 3 than one). This contention sits in tension with the idea that a belief might remain the same (have the
 4 same content, play the same functional role) whilst being sometimes conscious, sometimes not. In
 5 response, one might insist that the unconscious and conscious beliefs have different content - and
 6 perhaps sometimes they do; but to suppose they *always* did would be ad hoc.¹⁰
 7

8 Accordingly, another option might be to say that there is one belief (the racist belief) and
 9 that it is held both consciously and unconsciously. That is something we should want to avoid
 10 saying on pain of inconsistency. Alternatively, we might say that this belief is sometimes held
 11 consciously, sometimes unconsciously. This is surely true of many of our beliefs: my belief that
 12 walls are impassable is unconscious until I explicitly call to mind the properties (solidity, etc) that
 13 the wall has – yet it is causally efficacious even when it remains unconscious (I do not attempt to
 14 traverse the wall). This kind of move would help the unconscious beliefs view make good sense of
 15 non-conflict cases. But this then seems to require that we treat cases of conflicting biases quite
 16 differently from cases of non-conflict. In conflict cases (such as that of the biased egalitarian) the
 17 thought is precisely that appeal to an agent’s run-of-the-mill beliefs is insufficient to explain her
 18 behaviour – we need to consider other aspects of her cognition to explain what is going on.
 19 Mandelbaum suggests that we call that other aspect an ‘unconscious belief’, which conflicts with
 20 the agent’s conscious beliefs. However, in the non-conflict cases, we just appeal to the agent’s run-
 21 of-the-mill beliefs, which – like other beliefs – are sometimes conscious, sometimes not. The
 22 unconscious belief account then owes us an explanation of why the mechanism that produces the
 23 automatic and autonomic features of behaviour is different in conflict than in non-conflict cases.
 24

25 On the other hand, if it turns out that these ‘other aspects’ that explain the automatic and
 26 autonomic features of her discriminatory behaviour are just beliefs (sometimes unconscious ones, in
 27 the way that my belief about walls is sometimes unconscious, sometimes conscious and occurrent)
 28 then we seem to lose some explanatory power that we were supposed to gain by positing
 29 supposedly distinctively unconscious beliefs in the first place. Why is it difficult for agents to be
 30 aware of these racist unconscious beliefs if they are just run of the mill beliefs? (I don’t have this
 31 difficulty with my beliefs about the solidity of walls.) Perhaps there is an important distinction to be
 32 had between non-occurrent beliefs and unconscious beliefs – but we need to see what that
 33 distinction is and how it helps us to make sense of the role of implicit biases in non-conflict cases.
 34

35 Before we see how those claims might be fleshed out, we are not in a position to evaluate
 36 fully whether or not the unconscious belief account can satisfactorily meet our desiderata –
 37 especially in terms of providing a framework for moral evaluation; and for providing testable
 38 hypotheses for effective interventions to mitigate implicit biases (for example, how similar to
 39 conscious or non-occurrent beliefs do unconscious beliefs have to be for the same framework for
 40 moral evaluation to apply?). Thinking about our test cases shows us that there is more work to be
 41 done to complete the unconscious belief model and evaluate whether it is able to meet our
 42 desiderata.
 43
 44

45
 46 *(iv) Aliefs.* Gendler has recently introduced the notion of *alief* to capture certain aspects of our
 47 cognition that are different in kind from our beliefs and other propositional, reflective and rationally
 48 guided states. These other kinds of mental states, aliefs, are associative, automatic and arational. On
 49 Gendler’s characterisation:
 50

51
 52 A paradigmatic alief is a mental state with associatively linked content that is
 53 representational, affective and behavioral, and that is activated— consciously or
 54 nonconsciously—by features of the subject’s internal or ambient environment. Aliefs may be
 55

56
 57
 58 ¹⁰ One might insist that since they have different functional roles, they are different mental states: but on familiar views whereby
 59 mental states are individuated by their functional role, this will be different to reconcile with Mandelbaum’s insistence that the
 60 mental state at issue, albeit unconscious, remains a belief. Thanks to Komarine Romdenh-Romluc for this thought.

1
2 either occurrent or dispositional (2008a, p.642)
3

4 Aliefs are states constituted by tripartite clusters of co-activated contents: this includes
5 representational content, affective states, and the readying of motor responses. The representational
6 content of aliefs need not be propositional, and aliefs can be held consciously or non-consciously.
7 Positing aliefs, Gendler suggests, can help us make sense of the cognitions that are at work where
8 people display responses or dispositions that are at odds with their beliefs. I *believe* the Skywalk is
9 perfectly safe; but my *alief* is 'High up! [Representation]; Scary! [Affective response]; Get off!
10 [Readying of motor response]'.
11

12 We can readily think of implicit biases as a kind of alief, and indeed, the phenomenon of
13 implicit bias is one that Gendler uses to elucidate the notion of alief (2008b p.553). Consider our
14 example of the biased egalitarian (C): this individual has various egalitarian and anti-racist beliefs:
15 that the race of applicants is not relevant to the quality of the application, that her interactions
16 should not differ depending on the race of the interviewee. But her implicit bias, as manifested in
17 greater hostility in her interactions, can be unpacked as a cluster of activated contents: 'black
18 applicant [representational]; negative affect/fear [affective response]; hostile/avoid [readying of
19 motor response]'. This is the agent's alief towards the black applicant, which is discordant with her
20 other mental states, beliefs, that do not discriminate between the black and white applicants.
21

22 Gendler uses these discordant cases to elucidate the phenomenon of *alief* - cases in which
23 an agent's 'salient occurrent aliefs will come apart from her occurrent reality-reflective attitudes'
24 (2008b, p.554). But crucially for present purposes, we may also find aliefs that are norm-
25 concordant: cases in which an agent's occurrent aliefs are in accordance with her beliefs and other
26 explicit attitudes (2008b, p.554). These 'norm-concordant' cases are what we have in non-conflict
27 cases.
28

29 Accordingly, it looks like the alief model is well placed to help us to make sense of the test
30 cases that have, thus far, posed difficulties for models of implicit bias. We can say that the racist (A)
31 and protocol-adhering racist (B) have explicitly racist beliefs alongside the aliefs 'black applicant;
32 negative affect; hostile'. That is, these agents have a cluster of associative, automatic responses that
33 influence behaviour and judgement, alongside the role that their explicit mental states play. This
34 looks like a viable model for dealing with our difficult test cases, then. With respect to our other
35 desiderata, this model certainly departs from our existing theories of mind, and folk notions of
36 psychological state - and it will certainly take some careful work to articulate how aliefs might fit
37 into a framework for moral evaluation. If we take seriously Gendler's remarks that aliefs can be
38 changed, by processes of habituation or counter-conditioning of cognitive responses (2008b, p.572-
39 576), it looks like aliefs can provide a model that meets these desiderata. So, if the departure from
40 folk psychology is well motivated there may be reason to accept it.
41

42 However, I do not think it is well motivated, primarily for reasons expounded by Greg
43 Currie and Anna Ichino. They argue, first, that there is no reason to suppose that there are mental
44 states, aliefs, as such. They claim that the contents Gendler identifies as 'clustering' in a way that
45 underwrites a mental state - alief - may be causally related, but there is no reason to suppose they
46 constitute a 'substantial entity' or sui generis state. Indeed, secondly, they claim these contents are
47 best understood as familiar other states: perception or representation; affect and behavioural
48 response. Whilst Gendler emphasises the 'tightness' of the clustering, Currie and Ichino point out
49 that she also accepts that the causal relations between the contents can be broken - by habituation,
50 or practice, say (2012, p.790). This is well explained by a view that sees them as clusters of causally
51 related distinct states, rather than as one sui generis state, alief. In fact, they argue, many other states
52 are closely causally related (believing you are being charged by a bull, fearing it, readying to flee),
53 but there is no theoretical gain from identifying such a cluster with some new distinctive mental
54 state.
55

56 I find the case they make compelling, and do not see the motivation for gathering together
57 the automatic, associative states involved in implicit cognition into a new mental state, *alief*. In fact,
58 I think there is good reason to refrain from doing so, if this move gives the misleading impression
59
60

1
2 of homogeneity in the mental states we call ‘aliefs’. I have argued elsewhere (Holroyd & Sweetman
3 2016) that there is considerable heterogeneity in the sorts of associations involved in implicit biases,
4 and that recognising this is important for understanding the different strategies that may be needed
5 to combat implicit bias (see desiderata 5). For example, it may be that some of the phenomena we
6 call implicit bias are co-activated representations (such as those which encode stereotypes); others
7 co-activated representation and affect; others yet co-activated representations, affect and
8 behavioural response. We need not suppose that the same tripartite structure will be found in all
9 instances of implicit bias. Accordingly, we should not endorse a model of implicit cognition that is
10 unmotivated, and obscures important differences between different implicit biases.

11 Note that this line of enquiry has been fruitful, however: it points us towards an
12 understanding of implicit bias as constituted by different combinations of co-activated contents:
13 representational or affective or activated motor responses, which stand in close causal relations (this
14 may be the sort of ‘cousins’ of alief that Currie and Ichino suggest in their 2012, p.797). Perhaps
15 this minimalistic understanding is sufficient for modelling the phenomenon: further investigation
16 will be required to ascertain whether it satisfactorily meets our desiderata. But we can see how at
17 first glance, this ‘minimal model’ may find some empirical support. It coheres with understandings
18 of our psychology which recognise associative processing - though may depart from folk
19 psychological models that leave little room for this. Careful work will be needed to ascertain
20 whether this sort of cognition falls within the remit of moral evaluation or warrants ascriptions of
21 responsibility - but we may compare other sorts of automatic processing to make headway on this,
22 so have a framework to work with. And this kind of model can accommodate the diversity of
23 processes involved in implicit cognition - either representational, or affective, or behavioural -
24 which may be important in identifying the full range of phenomena that require intervention.
25 However, before it is worth pursuing such a model further, we should consider a challenge to the idea
26 that we need a model of implicit cognition at all.

27
28
29
30
31 *(v) Traits.* Perhaps our attempts to model implicit cognition are misguided, because in doing so we
32 suppose that there is such a thing as implicit biases or implicit attitudes. But on Machery’s view,
33 this distinction is itself wrong-headed, at least in relation to implicit attitudes (2016). Machery’s
34 claim is that there is no distinction between implicit and explicit attitudes, since attitudes are traits,
35 and so cannot manifest the property of being either implicit or explicit. The idea here is that, if we
36 take seriously the idea that attitudes are dispositional profiles, then we reject the claim that they are
37 mental states. Attitudes are dispositions to cognize, respond affectively, and behave in certain ways.
38 On this view, then, attitudes cannot be implicit or explicit, since they are traits, and traits are not the
39 sort of things that can be implicit or explicit.

40
41 What do we say, then, about those sorts of things we might have called ‘explicit attitudes’:
42 the attitudes of the biased egalitarian, for example, that seem to be at issue when she proclaims that
43 she affirms the value of racial equality? On Machery’s view, these are judgements about attitudes,
44 rather than expressions of attitudes themselves - judgements that we have certain attitudes. And,
45 these judgements can be more or less accurate: the biased egalitarian inaccurately thinks she is non-
46 discriminatory, but parts of the trait she is reporting on are not accurately detected; those
47 dispositions to discriminate in some contexts. Hence the departure between her report and the
48 dispositional profile that we observe when she discriminates.

49
50 Where does this leave the implicit/explicit distinction, such that we can make good sense of
51 our test cases above? Whilst I am generally sympathetic to the trait view of attitudes, it does not yet
52 provide the resources to help us make good sense of implicit biases per se. This analysis just pushes
53 the distinction back a level - attitudes themselves are not implicit; but components of attitudes may
54 be. On this view, implicit biases will be a component of the trait (along with various other things).
55 Some of the components (the psychological basis) of the trait may include the sorts of automatic
56 responses identified as part of implicit cognition, as Machery seems to acknowledge. So far as the
57 trait view offers us an theory of attitudes, it tells us that these attitudes are not, themselves, implicit
58 cognitions. But it does not give the rest of the story about implicit cognition. And insofar as there

1 are implicit biases or associations that constitute the psychological basis of attitudes we want to
 2 know what it is about *these* states that characterises them as implicit and delivers a model that meets
 3 our desiderata.
 4

6 IV

7 *Preliminary Conclusions:* I have suggested that five ways of thinking about implicit biases - as
 8 patchy endorsements, as in-between beliefs, as unconscious beliefs, as aliefs, or as part of the
 9 psychological basis of attitudes - are not yet satisfactory in helping us to meet our desiderata for an
 10 account of implicit cognition. They are unmotivated, or leave no room for, or obfuscate, or are
 11 simply yet unable to speak to, the explanatory significance of implicit biases in what I have called
 12 non-conflict cases - cases in which the agents implicit and explicit racial attitudes are aligned.
 13

14 I have suggested that what I referred to as the 'minimal model' above may fare better: this
 15 account sees implicit biases simply as causally related, or co-activated representational contents, or
 16 affective and behavioural responses. Some or all of these contents or responses may be involved in
 17 different cases in which implicit bias is operative. It seems to me fruitful to see what work the
 18 minimal model can do. There are two key thoughts that we should bear in mind when developing
 19 the minimal model. First, we should note that nothing in the minimal model entails that implicit
 20 biases are necessarily unconscious or impossible to be aware of. And I think there are good reasons
 21 avoid any such commitments. Recent empirical evidence suggests we may be able to gain
 22 awareness of the presence of implicit biases, or of their influence on behaviour. It might be difficult
 23 to have introspective access to biases, or to observe their effects in our cognitions and behaviours;
 24 but it is not clear that it is impossible to garner such awareness (see e.g. Pronin et al 2007, Hahn et
 25 al 2014).¹¹ That there may be variation in the extent to which awareness is possible in relation to
 26 implicit bias coheres with the minimal model, since we can draw on a host of explanatory resources
 27 to see why we might sometimes be alert to the activation of representational content or affective
 28 responses, and at other times unaware. This opens up the further question of in what sense this sort
 29 of cognition is 'implicit'.
 30

31 Secondly, we might explore, in relation to the minimal model, the various kinds of control
 32 that we have with respect to implicit biases, and whether these distinguish implicit from explicit
 33 cognitions. Whilst implicit biases are not beyond our control,¹² it seems to be the case that the kind
 34 of control we may have in relation to implicit biases differs significantly from the kind of control
 35 we can exercise over other states that are guided by reflective deliberation: a framework focused on
 36 distinctive kinds of control may be of use in developing an account of implicit cognition. And we
 37 now have articulated some desiderata, and some heretofore ignored test cases, that can assist us in
 38 developing a more satisfactory model.
 39
 40

41 Department of Philosophy
 42 The University of Sheffield
 43 45 Victoria Street
 44 S3 7QB
 45 j.d.holroyd@sheffield.ac.uk
 46
 47
 48
 49
 50

51 Thanks to audiences at Institutet för Framtidsstudier, Stockholm, the MAP conferences at KCL and
 52 Leeds, and to the CogSci research network at the University of Sheffield; each provided
 53

54
 55 ¹¹ I have discussed the different kinds of awareness we might have in relation to implicit bias in Holroyd 2014. I also
 56 consider various explanations for why we might be unaware of bias, even if such awareness is possible: self-
 57 deception, or confabulation, for example.

58 ¹² I have elsewhere argued, with Dan Kelly, for a model of ecological control that helps us to make sense of the kind of control at
 59 issue (Holroyd & Kelly 2015). See also Scaife (ms.) for a model of implicit cognition that focuses on the kinds of control we
 60 have over implicit biases.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

constructive discussion in which the ideas of the paper were developed. This research was supported by the Leverhulme Trust Research Project Grant, on 'Bias and Blame', and I have benefited from fruitful discussions with co-investigators Tom Stafford, Robin Scaife and Andreas Bunge. Special thanks to Anna Ichino and Federico Picinali, for invaluable discussions of earlier drafts.

References

Committee on Maximizing the Potential of Women in Academic Science and Engineering, National Academy of Sciences, National Academy of Engineering, and Institute of Medicine, 2007: 'Beyond Bias and Barriers: Fulfilling the Potential of Women in Academic Science and Engineering'. National Academies Press, <http://www.nap.edu/>

Cudd, Ann 2005: 'How to Explain Oppression: Criteria of Adequacy for Normative Explanatory Theories' *Philosophy of the Social Sciences*, 35(1), pp.20-49

Currie, Greg and Anna Ichino 2012: 'Aliefs don't exist, though some of their relatives do'. *Analysis* 72 (4), pp.788-798.

Patricia G. Devine, E. Ashby Plant, David M. Amodio, Eddie Harmon-Jones, and Stephanie L. Vance, 2002: 'The Regulation of Explicit and Implicit Race Bias: The Role of Motivations to Respond Without Prejudice'. *Journal of Personality and Social Psychology* 82, pp.835-48.

Devine, Patricia G. 1989: 'Stereotypes and Prejudice: their automatic and controlled components'. *Journal of Personality and Social Psychology* 56, pp.5-18.

Dovidio John F. & Samuel L. Gaertner 2000: 'Aversive Racism and Selection Decisions: 1989 and 1999'. *Psychological Science* 11, pp.319-323

Dovidio, John. F. Samuel, L. Gaertner, Kerry Kawakami, & Gordon Hodson 2002. 'Why can't we just get along? Interpersonal biases and interracial distrust' *Cultural Diversity and Ethnic Minority Psychology*, 8(2), 88-102.

Dovidio, John F.; Kawakami, Kerry; Johnson, Craig; Johnson, Brenda; Howard, Adaiah. 1997: 'On the Nature of Prejudice: Automatic and Controlled Processes'. *Journal of Experimental Social Psychology* 33(5), pp.510-40

Dovidio John F., Kawakami Kerry, Gaertner Samuel L. 2002: 'Implicit and explicit prejudice and interracial interaction'. *Journal of Personality and Social Psychology*; 82(1), pp. 62-8

Fischer, John Martin and Mark Ravizza 2000: *Responsibility and Control: A Theory of Moral Responsibility*, Cambridge University Press.

Gawronski, Bertram, Eva Walther, and Hartmut Blank 2005: 'Cognitive Consistency and the Formation of Interpersonal Attitudes: Cognitive Balance Affects the Encoding of Social Information'. *Journal of Experimental Social Psychology*, 41, pp.618-626.

Jack Glaser, Eric D. Knowles 2007: 'Implicit motivation to control prejudice'. *Journal of Experimental Social Psychology* 44, pp.164-172

Gendler, Tamar 2008a: 'Alief and Belief'. *Journal of Philosophy*, pp.634-663

----- 2008b: 'Alief in Action (and Reaction)'. *Mind and Language* pp.552-585

Glasgow, Joshua 2016: 'Alienation and Responsibility'. In Michael Brownstein & Jennifer Saul (eds.) *Philosophy and Implicit Bias*. Oxford University Press

Hahn, Adam, Charles M. Judd, Holen K. Hirsh, and Irene V. Blair 2014: 'Awareness of implicit attitudes'. *Journal of Experimental Psychology: General* 143 (3), pp.1369.

Hieronymi, Pamela 2008: 'Responsibility for Believing'. *Synthese* 161(3), pp.357-373.

Holroyd, Jules 2014: 'Implicit Bias, Awareness and Imperfect Cognitions'. *Consciousness and Cognition*, 33, pp.511-523

Holroyd, Jules & Daniel Kelly 2015: 'Implicit Bias, Character, and Control'. In Jonathan Webber and Alberto. Masala (eds.) *From Personality to Virtue*, Oxford: Oxford University Press.

Holroyd, Jules & Joseph Sweetman 2016: 'The Heterogeneity of Implicit Bias'. In Michael Brownstein & Jennifer Saul (eds.) *Philosophy and Implicit Bias*. Oxford University Press

Jost, John T., Laurie A. Rudman, Irene V. Blair, Dana R. Carney, Nilanjana Dasgupta, Jack Glaser, and Curtis D. Hardin 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 Organizational Behavior*, 29, pp.39-69

Lai, Calvin K., Maddalena Marini, Steven A. Lehr, Carlo Cerruti, Jiyun-Elizabeth L. Shin, Jennifer A. Joy-Gaba, Arnold K. Ho Arnold K.; Teachman, Bethany A.; Wojcik, Sean P.; Koleva, Spassena P.; Frazier, Rebecca S.; Heiphetz, Larisa; Chen, Eva E.; Turner, Rhiannon N.; Haidt, Jonathan; Kesebir, Selin; Hawkins, Carlee Beth; Schaefer, Hillary S.; Rubichi, Sandro; Sartori, Giuseppe; Dial, Christopher M.; Sriram, N.; Banaji, Mahzarin R.; Nosek, Brian A. 2014: 'Reducing implicit racial preferences: I. A comparative investigation of 17 interventions'. *Journal of Experimental Psychology* 143, pp.1765-1785

Levy, Neil. forthcoming: 'Implicit Bias and Moral Responsibility: Probing the Data'. *Philosophy and Phenomenological Research*

----- 2015: 'Neither fish nor fowl: Implicit attitudes as patchy endorsements'. *Noûs*, 49(4), pp.800-823

Lippert-Rasmussen, Kasper 2014: *Born Free and Equal? A Philosophical Enquiry into the Nature of Discrimination* Oxford University Press

Mandelbaum, Eric 2015: 'Attitude, Inference, Association: On the Propositional Structure of Implicit Bias'. *Noûs* doi: 10.1111/nous.12089

Machery, Edouard 2016: 'De-Freuding Implicit Attitudes'. In Michael Brownstein & Jennifer Saul (eds.) *Philosophy and Implicit Bias* Oxford University Press

McConnell, Allen R., and Jill M. Leibold 2001: 'Relations among the Implicit Association Test, discriminatory behavior, and explicit measures of racial attitudes'. *Journal of experimental social psychology* 37(5), pp.435-442.

1
2
3 Pronin, Emily, and Matthew B. Kugler 2007: 'Valuing thoughts, ignoring behavior: The
4 introspection illusion as a source of the bias blind spot'. *Journal of Experimental Social Psychology*
5 43 (4), pp.565-578.
6

7 Saul, Jennifer 2013: 'Unconscious Influences and Women in Philosophy'. In Fiona Jenkins &
8 Katherine Hutchison (eds.) *Women in Philosophy: What Needs to Change?* Oxford: Oxford
9 University Press pp.39-60
10

11 Schwitzgebel, Eric 2010: 'Acting Contrary to Our Professed Beliefs or The Gulf Between
12 Occurrent Judgement and Dispositional Belief'. *Pacific Philosophical Quarterly* 91, pp.531-553
13
14

15 Washington, Natalia & Daniel Kelly, 2016: 'Who's Responsible for This? Moral Responsibility,
16 Externalism and Knowledge about Implicit Bias'. In Michael Brownstein & Jennifer Saul (eds.)
17 *Philosophy and Implicit Bias* Oxford University Press
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60