

Doxastic wrongs, non-spurious generalisations and particularised beliefs

Cécile Fabre

I

The problem. Consider the following interactions, both of which involve someone whom I shall call Albert. Albert is white, was born a male and self-identifies as a man.

Posh Party: Albert is a guest at a posh party. He spots an African-American man (call him John) who is wearing a tuxedo, and forms the belief, purely on the basis of his skin's color and attire, that he is a waiter. In fact, John is the guest of honour.

Board Meeting: Albert has been invited as a consultant to speak to the Board of a FTSE-100 company. There is only one woman (call her Susan) in the Board Room. Albert forms the belief, purely on the basis of her gender, that she is an administrative assistant. In fact, Susan is a Board member.

Albert forms beliefs about John's and Susan's respective occupations with no evidence other than the context in which he sees them combined with a demographic characteristic which has long grounded structural injustices. His particularised beliefs are grounded in non-spurious generalisations: African-Americans are more likely to be waiters than guests at posh parties, and women are more likely to be administrative assistants than Board Members of FTSE-100 companies.¹

Some people take the view that Albert wrongs John and Susan merely by dint of holding those beliefs. This view is an application of the general thesis that holding beliefs can wrong the subjects of those beliefs irrespective of whether one manifests that belief. Call this the doxastic wrongs thesis (henceforth, DWT.)

We hold particularised beliefs about one another on the basis of demographic generalisations in a wide range of contexts, many of which seem morally unproblematic (Schauer 2006). For example, it seems morally unproblematic to assume that the young person standing outside a classroom is a student rather than a teacher or that the person speaking English with a French accent is French rather than Belgian. The challenge for DWT lies in identifying which particularised judgements based on non-spurious statistical inferences are morally problematic and which are not.

¹ I adapt the classic cases from Tamar Gendler (2011) and Sarah Moss (2018b, pp. 220-222) respectively.

Most of the cases discussed in the literature on doxastic wrongs involve beliefs which have the following features: they are held on the basis of statistically non-spurious generalisations about gender and ethnicity; they are false; they are negatively valenced (in societies such as ours, regrettably, being a waiter or an administrative assistant does not command as much esteem as being a posh party guest or a FTSE-100 Board member); they are held by members of socially privileged groups about members of socially under-privileged groups.

Now suppose that Albert's beliefs are true: John is a waiter and Susan is an administrative assistant. Does Albert wrong them and, if so, why? Or suppose that (in other cases) Albert's beliefs, though formed on the basis of non-spurious generalisations are false yet positively valenced, or true and positively valenced, and/or are held about a member of a socially privileged group. Does Albert wrong the subjects of his beliefs?

In this paper, I argue that a plausible Kantian argument for the view that Albert wrongs John and Susan in classic cases such as *Posh Party* and *Board Meeting* applies to non-classic cases in which the belief is true and/or positively valenced, and/or is held about socially privileged individuals. The paper proceeds as follows. Section II provides a brief account of doxastic wrongs. Section III rejects Mark Schroeder's argument to the effect that one cannot wrong someone by holding true beliefs about them (Schroeder 2018). Sections IV and V draw on a Kantian argument for DWT in respect of negatively valenced false beliefs and apply it to negatively valenced true beliefs. Sections VI-VIII address three objections and extend the scope of the Kantian view to the other non-classic cases. Section IX concludes.

Three final remarks. First, unless otherwise stated, when I speak of beliefs, I mean beliefs grounded in statistically non-spurious generalisations which track socially salient characteristics. I rely on Kasper Lippert-Rasmussen (2011, p. 30)'s definition of a socially salient characteristic as a characteristic of a group such that 'perceived membership of it is important to the structure of social interactions across a wide range of social contexts'. Second, I focus on beliefs, not credences. I remain agnostic on whether having a credence in p can wrong the subject of that credence. I thus assume that beliefs are not reducible to credences, not least because as Laura Buchak (2014) argues, some of our evaluative practices, such as blame attribution, rely on beliefs, not credences. (Buchak focuses on blame. I conjecture that her argument extends to at least some of the evaluative judgments I discuss in this paper.) Third, I neither attempt to provide a full argument for DWT in general nor suggest that the Kantian argument I offer in its defence is the only way to vindicate it. My more modest aim is to show that if DWT holds in the classic cases on the aforementioned Kantian grounds, it also holds in the non-classic cases on those very same grounds.

II

Doxastic wrongs: a primer. Suppose that some agent X believes that p about some other agent Y . DWT says that X can thereby wrong Y . The wrong is located in believing that p - as distinct from acting on that belief - and is directed at the subject of that belief

(Basu and Schroeder 2018).²

Consider the location of the wrong. DWT avers that Albert wrongs John by believing that he is a waiter - irrespective of his behaviour. The wrong is located in the way in which Albert forms his belief, or in his mere holding the belief irrespective of how he came to form it, or both. In our opening example, he forms that belief based on perceptual evidence together with his knowledge of the relevant demographics, yet discounts the fact that John might not in fact *be* a waiter (a point to which I shall return in s. IV.) However, having formed a (relevant) belief is not a necessary condition for committing a doxastic wrong. Suppose that Albert swallows a pill one effect of which is that he instantaneously holds stereotypical beliefs about any African-American male he encounters. Let us concede that he has not *formed* the belief that John is a waiter. Suppose next that he realises that not everyone holds stereotypical beliefs about African-Americans and that there is evidence out there which undermines those stereotypes - stereotypes, he also knows, which are morally controversial. He makes no effort to investigate his pill-induced belief in the light of this evidence. It is intuitively plausible that he wrongs John notwithstanding the fact that he has not formed the belief that John is a waiter. His wrong is located in his holding and refusing to try and school himself out of that belief.

Consider now the direction of the wrong. On some views, Albert's wrong lies in his failure to display doxastic virtues or to live by the principle that, in general, we morally ought to keep an open mind about people. DWT goes further: by assuming, on the basis of John's skin color together with the context in which he encounters him that John is a waiter, Albert wrongs *him*.

DWT is a tough sell - implying as it does that there are such things as (in common parlance) 'thought crimes'. Yet suppose that John overhears Albert refer to him as a waiter. When challenged by John, Albert apologises for not finding an isolated spot before placing the call. His apology misfires. The problem is not just that he has unnecessarily run the risk of offending John. The problem is that he has formed a false belief about John with no evidence other than a context- and race-based generalisation. *That* is John's grievance, and it is not addressed by Albert's apology. Moreover, Albert owes an apology *to John*: were he to say to a fellow guest that he is sorry for mistakenly assuming that John is a waiter, he would also miss the mark.

DWT elicits two standard objections. First, according to the lack-of-control objection, we are under a moral obligation not to φ only if we are morally responsible for φ -ing; we are morally responsible for φ -ing only if we have some degree of control over φ ; given that we do not have control over what we believe, there is no such thing as a doxastic moral obligation and, therefore, as a doxastic wrong.

² DWT thus differs from one of the earliest accounts of the wrongfulness of badly formed beliefs, as articulated by William K. Clifford (1877) in his seminal paper 'The ethics of belief'. Clifford argues that, irrespective of the truth value of one's belief, believing on the basis of flawed evidence is morally wrong. However, he does not explicitly state that the agent wrongs the subject of his beliefs; nor does he seem willing to locate the wrongfulness of the belief in the belief *per se*, rather than on the consequences of holding the belief on our actions.

The objection's normative premise, that having control over whether to φ is a necessary condition for being under a moral obligation not to φ , has come under sustained criticism, including in the context of belief (Smith 2005; Hieronymi 2006, 2008; Basu and Schroeder 2018). But even if the normative premise is true, its factual premise, that we lack control over what we believe, is not. Many of our practices are intelligible only if we take ourselves to have such control and if, at least sometimes, we are proved to be right. We undergo cognitive behaviour therapy in the hope, often fulfilled, that we will be led to revise noxious beliefs. We embark on scientific inquiry in the hope, often fulfilled, that we will acquire justified beliefs and shed unjustified ones. In many jurisdictions, the criminal law presupposes that we have some degree of control over some of our mental states in general, and beliefs in particular. The costs of rejecting that view include, for example, not penalising a doctor who believes that his patient's ventilator is functioning when the evidence clearly points to it being faulty. When challenged on their manifest prejudices, some individuals have been found to respond to implicit attitudes tests in less prejudiced ways (Basu 2019a, pp. 2513-2514; Basu 2018; Greenberg 2020; Holroyd 2012, pp. 282-286; Marusic and White 2018). All that is needed for DWT to hold is that we have and can acquire some degree of effective control over some of our beliefs.

The second objection goes like this. Reasons for holding and rejecting beliefs can only be epistemic reasons - grounded in whether the beliefs are true, what evidence there is to support or undermine them, and so on. The fact that a belief has morally problematic features is neither here nor there.

In response, some proponents of DWT argue that moral considerations can make a difference to the epistemic status of a belief. In now accepted parlance, the former *encroach* on the latter. Moral encroachment is a variant of pragmatic encroachment, according to which the epistemic status of a belief in part depends on the believer's pragmatic circumstances. According to proponents of moral encroachment, by parity of reasoning, the epistemic status of a belief can also depend on moral considerations in favour or against holding the belief (Pace 2011; Fritz 2017; Basu and Schroeder 2018; Moss 2018a; Schroeder 2018).

Moral encroachment is controversial. In any event, one need not endorse it in order to defend the doxastic wrong thesis from the objection under scrutiny. One can grant that a belief is epistemically rational yet morally wrong (Fritz 2020). Alternatively, one can focus solely on epistemic norms. If and when the believer wrongs the subject of his belief by relying on patchy evidence when he can reasonably be expected to have or procure better evidence, the believer's moral failure is traceable to his epistemic failures (Begby 2018; Begby 2021, pp. 165-170).

In what follows, I remain agnostic on (a) the soundness of moral encroachment as an account of the relationship between the moral and epistemic features of beliefs, and (b) the relationship between moral encroachment and DWT. All I need is the view, for which I have provided some intuitive support, that one can sometimes wrong another person merely by holding some belief about them. The question is which beliefs are of that sort.

III

The false-beliefs restriction. In the classic cases, Albert's beliefs are false: John is a guest, and Susan is a Board member. Let us re-label the cases *Posh Party_F* and *Board Meeting_F*. Mark Schroeder, who endorses the doxastic wrong thesis, restricts its scope to cases such as these: if Albert's beliefs are true (*Posh Party_T* and *Board Meeting_T*), he does not wrong John and Susan.

On Schroeder's account, what makes holding a belief wrongful to its subject is that being erroneous is costly even if we do not manifest that belief. This is because of the place occupied by our beliefs in our relationships with other people. As he puts it (Schroeder 2018, p. 121),

our interpersonal relationships are in part constituted by our beliefs about one another. In so far as our beliefs help to constitute our relationships, the effects of our beliefs on our relationships are not mediated by the effects of our beliefs on our actions or other behaviours.

I agree that relationships are damaged when the beliefs which help constitute them are false beliefs. By implication, though, this argument does not provide the right explanation for cases in which the holder and subject of the belief fleetingly encounter each other, as in *Posh Party_F* and *Board Meeting_F*. Moreover, it does not show that the doxastic wrong thesis only applies to false beliefs: identifying one particular way in which holding certain beliefs damages relationships does not rule out other ways in which it can do so.

When handling cases in which the belief is true, Schroeder makes two moves. First, in some cases, one's true particularised beliefs about a person depend on false generalisations about the class to which she belongs. Suppose that Albert forms the (correct) belief that Susan is an administrative assistant and not a Board Member on the grounds that women generally do not have what it takes to rise to the top. Even though there is one respect in which Albert wrongs Susan in *Board Meeting_F* but not in *Board Meeting_T*, 'there *may* be some further respect' in which he wrongs her in both cases, 'as [a member] of a class' about which he holds false essentialising beliefs (p. 122). Given that his particularised belief, though true, relies on false beliefs, we can account for the view that he wrongs Susan without extending the thesis to true beliefs.

Schroeder is correct that it can be directedly wrongful to hold beliefs which are false with respect to their particular subject, and beliefs which are true with respect to their particular subject but which are grounded in generalised false beliefs. However, this move alone does not rule out adding true beliefs grounded in non-spurious generalisations to the category of doxastic wrongs.

Schroeder's second move goes like this. We must distinguish between objective and subjective moral permissibility. Suppose that my neighbour tells me that her daughter Jane, who likes cycling in our shared driveway, is down for a nap and will not wake up until 3pm. As I am typing this sentence sitting at my desk at 1:45pm, my neighbour's testimony is good enough for me to believe that Jane is safely asleep bed. At 2pm, I decide to go for a drive. My belief happens to be true. As a matter of practical ethics, I

am objectively permitted to drive off without checking and thus do not wrong her by doing so, but my failure to carry out basic visual checks means that I am not subjectively permitted to drive off. The distinction between subjective and objective permissibility applies to beliefs, as well as to actions. Thus, even if Albert does not wrong Susan - for he accurately and objectively believes that she is an administrative assistant - in some cases his belief might 'nevertheless be subjectively wrong, epistemically irrational, and therefore incapable of being knowledge' (Schroeder 2018, p. 123). This, Schroeder says, helps accounts for the intuition that there is something morally problematic about holding these kinds of belief - though Albert falls short of wronging Susan.

I agree that if Albert forms his belief about Susan on the basis of faulty evidence and/or as a result of sexist biases, it is plausible to say that he is not subjectively permitted to hold that belief. However, the distinction between subjective and objective permissibility cuts across the distinction between an act or a belief being wrong *simpliciter* and an act or a belief wronging its subject. Schroeder seems to rule out the possibility that Albert subjectively (albeit not objectively) wrongs Susan. Not only does this move seem unwarranted: there are independent reasons to resist it. For it seems to imply that, as a matter of practical ethics, I do not subjectively wrong Jane by driving off without checking whether she is cycling in the driveway. Yet, on plausible accounts of the morality of risk imposition, my failure to take her vital interests into account in my deliberation is a dereliction of a duty which I owe *her*. By parity of reasoning, Albert's failure to take due care in forming his belief about Susan wrongs her even though his belief happens to be true.³

In the next two sections, I develop this point into an argument for extending the thesis to cases such as *Posh Party_T* and *Board Meeting_T* and apply it to cases in which the subject of the belief is a wrongdoer.

IV

The Kantian View. Albert forms his true belief about John and Susan on the basis of a statistically accurate generalised belief: it really is the case that African-Americans are more likely to be waiters than guests at posh parties and that women are more likely to be administrative assistants than FTSE-100 Board Members - facts of which he is well aware. Still, it does not follow from the fact that John is African-American and from the fact that any African-American is more likely to be a waiter than a guest at a posh party, that *John is* a waiter. Albert only has epistemic warrant for the claim that John is more likely to be a waiter than a guest. Likewise, *mutatis mutandis*, regarding Susan. The question is whether, and if so why, he wrongs John and Susan by accurately believing that they are, respectively, a waiter and an administrative assistant.

Rima Basu's Kantian defence of DWT, though focused on cases in which the belief

³ I draw on John Oberdiek's recent account of the morality of pure risk imposition (Oberdiek 2017, esp. ch. 4). In the course of his sustained critique of Schroeder's view as applied to false-belief cases, James Fritz (2020) avers without further ado that the epistemic status of a belief as true or false in itself has little to do with its being morally suspect.

is false, is a useful starting point (Basu 2019b, p. 922ff).⁴ Let us accept the Kantian principle that we ought to relate to one another as persons, and not as objects. For me to treat you as a person requires that, in my dealings with you, I recognise that your occupation, your relationships, your political and religious views and your hobbies are constitutive of your identity and your sense of your own worth. To treat you as a person means that I owe it to you that my beliefs about you in those dimensions should be responsive to and match the ways in which you describe and understand yourself. Failing that, I do not relate to you as a person: rather, I form views about you as if I were ‘observing a planet’ whose movements are only answerable to the laws of physics.⁵

The Kantian View is demanding (a point to which I shall return in s. VI.) It also needs revising. Suppose that Albert’s belief is true, and that John and Susan identify with their job, or that they actively dis-identify with it, or that they could not care less about it. In the first case, Albert’s beliefs match John’s and Susans’ self-description. In the second and third cases, their self-description is not engaged by his beliefs: those beliefs are about some aspect of theirs (their job) which is not constitutive of their self-worth. Yet, intuitively, it still seems that Albert wrongs John and Susan in all of those cases as well.

Suitably revised, the Kantian View shows why Albert wrongs John and Susan. The key injunction is not that our beliefs about others should match their own views about who they are and their sense of their own worth. Rather, to treat others as persons is to recognise in them the capacity and willingness to act autonomously, that is to say, to frame, revise and pursue a conception of the good with which they identify. As Rawls (1993, pp. 48-54) would have it, it is to recognise that they have the moral power of rationality. This in turn requires that, when forming and holding beliefs about them, we owe it to them to be sensitive to whatever evidence we have that they are conducting their lives as autonomous agents. As the point is sometimes put in the literature on the wrongfulness of certain forms of discrimination, we owe it to them to treat them as individuals (Eidelson 2013, 2015). *Absent* such evidence, we owe it to them to give them the benefit of the doubt. Albert owes it to John and Susan to ascertain what they do before reaching a verdictive judgement about their occupation. If he cannot, will not or indeed may not do so (lest, for example, he should unduly intrude on their privacy), he owes it to them to allow for the possibility that they do not fit the stereotype. In doing so, crucially, he need not occlude the structural constraints under

⁴ In that paper, Basu, who draws on Peter Strawson (1962)’s interpretation of the Kantian requirement to treat one another as persons, explicitly says that one can wrong someone by holding true beliefs about them, but she does not fully explore that possibility: indeed, her discussion proceeds almost entirely with reference to cases in which the belief is false. See p. 920 for her explicit statement, and p. 919 for the exception, involving a racist hermit who holds the belief that the Indian person he sees on a photo must smell of curry - a belief which, at the precise time at which it is held, happens to be true. This case involves a particularised belief grounded in a spurious stereotype and thus falls outside the scope of my inquiry here. I much agree with Lawrence Blum (2004)’s discussion of the moral wrongness of this kind of stereotypes.

⁵ Marusic and White (2018) offer a similar defence of doxastic wrongs in false beliefs cases. The phrase ‘observing a planet’ is Rae Langton’s (Langton 1992, p. 486).

which they operate; nor need he eschew all probabilistic beliefs, based on statistically non-spurious and relevant generalisations, about them. If he is asked by his dinner companion whether it is more likely than not that John is a waiter and answers in the affirmative, he does not wrong John. If he is asked whether John *is* a waiter and so answers, he does. The important and familiar point is that he should not essentialise John by reducing him to the fact that he is African-American. Again, likewise *mutatis mutandis* with respect to Susan.

The claim that we are under a duty to others to give them the benefit of this particular doubt mirrors in the domain of beliefs precautionary principles for risk imposition under conditions of uncertainty (Bolinger 2020, pp. 2422-2426). So long as one grants that holding certain beliefs on the basis of non-spurious generalisations can wrong the subjects of those beliefs and that we owe it to others to take precautionary steps so as not to wrong them when faced with uncertainty about the relevant facts, one can endorse the claim that Albert owes it to John and Susan not to assume that they lack autonomy. The point holds irrespective of the truth-value of this assumption.

V

Wrongdoers. The Kantian View accounts for the claim that Albert commits a doxastic wrong in both variants of *Posh Party* and *Board Meeting*. In those classic examples, Albert's belief is negatively valenced: being a waiter or an administrative assistant does not command as much esteem as being a posh party guest or a FTSE-100 Board member. The Kantian View, if successful in cases in which the belief is both true and negatively valenced in that sense, also applies when the belief is true and negatively valenced in the stronger sense that its subject has committed a grievous wrong.⁶

Consider:

Knife Crime: Albert is taking his usual walk through the London Borough of Lambeth. He spots a young Black teenager (call him James) crouching next to the lifeless body of another Black teenager. Police and ambulance sirens can be heard in the distance. Albert forms the belief that James stabbed the lifeless teenager, in yet another episode of gang warfare.

According to figures released in 2018 by the office of the Mayor of London, slightly over four fifths of the knife crimes which were committed in London in 2017 and which resulted in injuries were committed by men, and slightly over two thirds by members of ethnic minorities. Two thirds of the victims of non-domestic knife crimes are BAME and almost all are male.⁷ Year on year, the Boroughs of Lambeth and Southwark witness

⁶ In his important article on the morality of relying on statistical evidence in criminal and civil trials, David Wasserman (1991) deploys the autonomy argument against such reliance; however, he restricts it explicitly to cases in which the defendant's liability is not conceded.

⁷ The figures are available at https://www.london.gov.uk/sites/default/files/mopac_justice_matters_13_march_2018_disproportionality_slides.pdf (accessed on 30th November 2021).

the highest number of knife crimes in the capital. Albert's particularised belief about James does not rely on a spurious generalisation - yet there is something deeply morally problematic about it.

The requirement that wrongdoers should not be treated on the basis of their race seems uncontroversial enough. The point here is that they ought not to be *thought of* on that basis either. This seems clear if Albert's belief is false (as when James is a friend of the murdered victim and would have come to harm too had he arrived on the scene ten minutes earlier.) But it is so, too, if Albert's belief is correct. For he holds the belief that James killed the teenager on grounds relevantly similar to those on the basis of which he formed accurate beliefs about John in *Posh Party*_T. Moreover, he does James an additional wrong, compared to his doxastic treatment of John. Not only does he assume of James either that he has not escaped the constraints of his socially salient characteristics or that he is not conducting his life in reflective awareness of those constraints - in denial of James' capacity for autonomy. In addition, he fails to respect him as a moral agent. To treat others as persons is not merely to recognise in them the capacity and willingness to act autonomously. It is also to recognise in them the capacity for moral agency and willingness to live accordingly - that is to say, the capacity and willingness to form considered judgements about right and wrong, and to lead their (autonomously chosen) life in reflective awareness of those moral judgements. In Rawls (1993), this is the second of our two moral powers - the moral power of reasonableness. It is far less extensively discussed in liberal political and moral philosophy than the power of rationality. Yet it is central to our moral lives. In the present context, Albert reduces James to the combination of three of his 'crime-relevant' demographic characteristics (age, gender and ethnicity). While he may not have evidence *that* James is abiding by the moral prohibition on murder, he also lacks evidence to the contrary. He owes it to him to assume that he may have formed the belief that killing is morally impermissible under the circumstances and lived his life accordingly.⁸

To say that Albert wrongs James by not suspending judgement is to imply that James has a grievance against Albert, and that Albert owes James an apology for believing that James killed the boy. But (it may be objected), in the light of his considerably more serious wrongdoing, how can James justifiably make such a demand of Albert, whose belief as to his culpability, let us not forget, was accurate?

The question is not specific to doxastic wrongs. It applies to practical wrongs and is raised by the racially-driven imposition of hard treatment. If Albert openly berates James for committing murder (a belief he forms on the basis of James' demographic characteristics), he acts without warrant even if James did in fact commit murder. Albert owes James an apology, it seems, even though James' wrongdoing is far worse than

⁸ Readers will have noticed that in this and the previous section, I speak of the capacity *and willingness* for rational and moral agency. Strictly speaking, however, in liberal moral and political philosophy, the focus is on capacity, not willingness. Yet it seems to me that they go hand in hand. Thus, it is a failure of respect on the one hand to say of another person that they are able to act morally (and thus are able to differentiate between right and wrong and act accordingly) yet on the other hand to assume without evidence that they are unwilling to do so (and thus did commit murder.)

Albert's. There is no reason not to apply the point to doxastic wrongs. Granted, in both cases, and particularly if James does not express remorse, we may think it is wrong of him to demand an apology. But the fact remains that Albert owes it to him to issue it.⁹

The Kantian View seems vulnerable to the following three objections: it is too demanding; it rides roughshod over the fact that, in some cases, prejudicial beliefs are a basis for morally permissible conduct; it proves too much.¹⁰I address them in turn.

VI

The demandingness objection. Consider the demandingness objection. Even though we have some degree of control over our beliefs, we often have unbidden thoughts about other people. Not only can we not help doing so and, in that respect, relying in one way or another on generalisations: it also saves us a huge amount of cognitive labour, thanks to which we can go about our lives.¹¹ Even if we have some degree of control over whether to believe p or $not-p$, it is very difficult for us to switch from believing p to believing $not-p$. It requires that we spend considerable amounts of time and cognitive energy examining relevant evidence against our existing beliefs. It also requires that we shed our beliefs if necessary, at the often painful cost of revisiting the relationships and experiences on the basis of which we formed them. Given the costs of the requisite epistemic labour, it seems harsh to say that we wrong those other people merely

⁹ Although I focus on particularised beliefs grounded in non-spurious generalisations, it is worth noting that the Kantian View also applies to wholly eccentric particularised beliefs. Suppose that Albert believes that red-haired men are more likely to be pastry chefs than blonde-haired men, and that short-haired women are more likely to commit murder than long-haired women. When encountering red-haired men and short-haired women in scenarios relevantly similar to *Posh Party* and *Knife Crime*, he forms a particularised judgement about those individuals on the basis of his eccentric generalisations. The Kantian View seems committed to charging Albert with wrongdoing. To some readers, this seems implausible: his beliefs are so strange that it is hard to imagine that their subjects could have a genuine grievance. I wonder, though, whether what drives this particular intuition is the thought that Albert cannot possibly be held epistemically responsible, and by implication morally responsible, for this belief, precisely because it is so strange. If he cannot, then he commits no wrong. (Thanks to Bill Winge for both the point and the pastry chef/hair colour, and to Alexander Greenberg and Guy Longworth for helpful oral and written exchanges on this.)

¹⁰ I am very grateful to the commentators and audiences mentioned in the Acknowledgements for raising many of the problems I discuss in this section.

¹¹ There is a very large literature on this issue, which Endre Begby (2021) nicely summarises in the context of a book-length account of the epistemology of prejudice. Two points. First, I set aside the complex question, at the heart of Begby's account, of whether someone who grows up in a deeply racist society is epistemically justified in holding racist beliefs. (See also Amia Srinivasan (2020)'s recent discussion.) Second, as Jessie Munton (2021) shows, one way in which we save cognitive labour is by ordering our existing beliefs, or failing to acquire new beliefs, in ways which give greater salience to some demographic features than to others. Salience however can manifest prejudice, as when (Munton's example) a male academic unthinkingly discards works by female scholars. By the lights of the Kantian View, we owe it to another not to give undue salience to some of their demographic characteristics, even if it turns out that our particularised beliefs about them is accurate.

by holding those beliefs - and even harsher still, to the point of implausibility, when those beliefs are true. Suppose that Albert has been brought up in a traditional, father-breadwinner/mother-at home family. This fact together with the context in which he encounters Susan explain why he assumes that she is an administrative assistant. To say that he wrongs her *merely* for holding that belief and *a fortiori* if he is not aware of being improperly biased, is unduly demanding of him.

Reply: of course there is a limit to what we may reasonably demand of one another; and the constraints of our professional, social and familial lives, the failures of our will and imagination, our upbringing, and the contexts in which we conduct ourselves clearly do make it difficult for us to revise our beliefs. (Brownstein and Saul 2016a, 2016b). But once we become aware of those beliefs, it is harder to exculpate ourselves for holding them. In, say, 1950 Britain, Albert could be excused for not being aware of the notion and risks of implicit biases: not so in 21st Britain (Holroyd 2012; Saul 2013; Washington and Kelly 2016; Zheng 2016). Moreover, even if Albert has an excuse of holding these beliefs here and now, it does not follow that he may not be called to account for holding them. Generally, one can wrong another person by φ -ing even if one is blameless for φ -ing - as proponents of the view that there is such a thing as strict moral liability would aver (Capes 2019). It is appropriate to call Albert to account, just as it is appropriate to call to account someone who makes objectively racist comments without properly realising what it is that he is saying (Calhoun 1989). One way to call him to account is to insist that he should at least try to revise his beliefs. Should he fail to do so, his (blameworthy) failure adds to his initial (blameless) wrongdoing to Susan.¹²

VII

The permissible conduct objection. A second putative difficulty with the Kantian View is this. In some cases, holding beliefs on the basis of generalised judgements about others is not merely understandable for the aforementioned reasons: it provides a basis for a rational and morally permissible response to one's predicament. Consider a woman - call her Betty - who is often sexually harassed by men while travelling alone on public transports or walking down the street. As a result, whenever she sees a man boarding her railway carriage or walking towards her, she cannot help believing that he will end up harassing her, and moves to a different carriage or crosses the street. Under those circumstances, it is rational and morally warranted (objectively speaking) of her to do that. Do we really want to say that she wrongs those men, not merely when her belief happens to be false but even when it happens to be true, given that it is traceable to repeated threatening encounters *and* that it leads her to protect herself from an actual threat? Surely not.

Reply: the claim that Betty is morally justified in protecting herself from this kind of

¹² Thanks to Roger Crisp for the point about strict liability. See also Khaitan and Steel (2018). I am very grateful to Guy Longworth for helping me to clarify my thoughts in this paragraph.

threat does not imply that the steps she takes to do so are exempt from moral evaluation in general, and from the charge of wronging their targets in particular. The objection goes through only if Betty's holding the relevant belief, on the basis of which she takes rational and warranted protective measures, does not wrong those men. This however is precisely the point at issue: so stated, the objection begs the question. Pending further objection and in the light of the Kantian View as applied to wrongdoers, we can still say that Betty wrongs those men even though her belief is correct. However - and this is crucial - she is not blameworthy, precisely because of the recurrent wrongdoings which led her to form it in the first instance. There is no inconsistency there: one can be under a duty not to φ and yet not be blameworthy for φ -ing.

Note that my point is compatible with the claim that an agent is sometimes justified, even obliged, to take protective action for the sake of another person yet without needing to form any belief at all, be such belief probabilistic or verdictive, about the specific threat which that person is facing. Suppose that Albert is tasked with protecting Vice-President Kamala Harris and notices a swastika-tattooed White male coming very close to her, gesticulating, during an impromptu downtown walk-in. Common-sensically, Andrew is under a moral obligation to act as if the man is more likely to be a threat to the VP than the middle-aged African-American woman standing at a similar distance bearing flowers. But he need not actually believe that the man is likely to be more of a threat than the woman. He can be agnostic yet morally required not to act on his agnosticism. That being said, an agent who is and believes herself to be under strong protective duties to another person, particularly a close relative or friend, is likely to find it harder not merely to resist from acting on prejudicial beliefs but to resist forming, or to revise, those beliefs. This takes us back to the problem of demandingness: what I said then applies here too.¹³

VIII

The proves-too-much objection. So far, my examples have involved beliefs which, whether true or false, are negatively valenced and are held by members of privileged groups about members of disadvantaged groups. A third difficulty with the Kantian View is that, by its own light, it implausibly charges with doxastic wrongdoing agents whose beliefs are about members of privileged groups and/or are positively valenced. Here are some cases, with matching replies.

Consider:

*Board Meeting_T**: Albert has been invited as a consultant to speak to the Board of a FTSE-100 company. There is a white man wearing a suit, in his mid thirties, standing near the head of the table (call him William).

¹³ As Guy Longworth pointed out to me, my point implies that to hold a 'silent belief' (on which one does not act) can be worse than to act as if one held that belief while not, in fact, holding it. I am thankful to Matthew Clark and Chase Mizzell for their help in thinking through the problem of protective duties and permissions. (Clark's example, in our oral exchange, involved Martin Luther King. Mine is reflective of contemporary America and combines gender-hatred with race-hatred.)

Albert forms the belief on the basis of his race and gender that he is a Board member. He happens to be correct.

On the Kantian View, to the extent that Albert fails to consider the possibility that William's position is not merely a by-product of his race and gender, he fails to give him the respect he is due as an autonomous agent. This is so whether or not William's conception of himself and the importance he attaches to his position match Albert's beliefs.¹⁴ Yet it seems implausible (the objection goes) to say that Albert *wrongs* William: to be thought of as a FTSE-100 company Board Member is rather a good thing, particularly when the belief is correct. In fact, to the extent that Albert's belief about William is also, by implication, a belief about women (that they are not Board members), it is those women whom Albert wrongs, not William.

Reply: the fact that Albert wrongs those women is compatible with the claim that he wrongs Williams. Moreover, the fact that Albert's belief is in one respect a good thing for William does not entail that it is not wrongful to him. Suppose that Albert is a local celebrity who has been asked to handout cash prizes at a showjumping competition. He is late and misses the show. The officials point out to him where the top three competitors stand: two of the competitors are female, and one is male. Without bothering to check the results, he hands out the top prize to the male rider, automatically assuming that he won on the grounds that (he thinks) male riders are more likely to win those sorts of events than female riders. Let us suppose that the generalisation on which his particularised belief rests and the particularised judgement itself are correct. Although there is a sense in which he benefits the male rider, there is also a sense in which he wrongs him, by failing to give him the reward for the right kind of reason - namely that he was the most skilled rider on the day. (The fact that *this* is the right kind of reason is compatible with the view that his being the most skilled rider on the day can itself be explained in part by structural gender-based injustices.) This case is relevantly similar to *Board Meeting_T**. If giving a reward on the aforementioned grounds does not exonerate Albert from wronging the male rider, forming a positively valenced belief on the aforementioned grounds does not exonerate him from wronging William.

Consider next:

Beaten Up_T: Albert spots a young African-American teenager who has clearly been beaten up. A white police officer is leaning over the teenager. Albert assumes on the basis of both parties' race and against a background of growing awareness of police brutality against African-American men that the police officer - call him Warren - is responsible for the beating. He happens to be correct.

It is implausible to say that Albert wrongs Warren by correctly believing, on the aforementioned basis, that he is a racist thug who abuses his official position. Warren is not the victim here. The teenager is.

¹⁴ In framing the Kantian View in this way, I part company with Basu's treatment of doxastic wrongs and members of privileged groups (Basu 2019b, p. 929).

Reply: if Albert wrongs James in *Knife Crime* by assuming on the basis of race-based stereotypes that he is responsible for killing the teenager, he wrongs Warren by making exactly the same kind of assumption based on relevantly similar stereotypes. This is clearly compatible with averring in both cases that the murdered teenager is a, indeed the greater, victim.

Finally, consider:

Posh Party Reverse_T: John notices Albert walking up the steps. He assumes, based on Albert's race and the fact that he is a guest, that Albert in turn will assume that he is a waiter on the basis of *his* race and the context of their encounter. He happens to be correct.

Surely - the objection holds - John is the victim here. It seems wholly implausible to say that *he* owes Albert an apology for believing, presumably on a long standing experience of interacting with posh white men, and as it happens correctly, that Albert will be guilty of stereotyping him.

Reply: the fact that Albert wrongs John does not imply that John does not wrong Albert. By the lights of the Kantian View, they wrong one another.

The fact that the Kantian View charges Albert with wronging the subjects of his beliefs in those cases is no reason to reject it. There is an outstanding worry, however. The Kantian View seemingly treats on a par *Board Meeting_T* and *Board Meeting_T**, *Knife Crime_T*, and *Beaten Up_T*, *Posh Party_T* and *Posh Party Reverse_T*. Yet intuitively, Albert commits a greater wrong in the first case of each pair than in the second. Can we account for this intuition while remaining faithful to the Kantian View?

This is a familiar difficulty in the literature on discrimination in general and racial profiling in particular. Profiling which targets members of ethnic minorities and employment practices which discriminate against women (for example, in the latter case, in respect of bodily strength for certain jobs) are particularly problematic in so far as they are parasitic upon and further entrench existing injustices, even if the statistical generalisations on which those practices rest are not spurious and are not intended so to discriminate (Mogensen 2019; Eidelson 2015, ch. 6, ch. 6; Lippert-Rasmussen 2013, ch. 11).

However, *ex hypothesi*, Albert does not act on or manifest his beliefs, and the claim that he further entrenches existing injustices seems of little help here. Could it be, then, that the reason why his doxastic wrongs are worse in the first case of each pair is that, as Basu (2019a, p. 924) puts it, John, Susan and James are 'more dependent than [members of privileged groups] on external validation for the maintenance of self-respect and self-esteem', precisely because of the structural injustices to which they are routinely subject? That they are 'dispositionally vulnerable' to failures of respect, even if the beliefs at issue are correct?

This is an intuitively appealing thought. But we still need to know why they are vulnerable to 'silent' beliefs. The answer lies in the ethics of risk imposition. In all but the most rarefied cases such as that of a hermit, we cannot but communicate the beliefs we hold and the reasons for our holding them; and, more often than not, those beliefs manifest themselves, sooner or later, in various forms of treatment. When particularised

beliefs based on socially salient characteristics are negatively valenced, are false and take members of disadvantaged groups as their subjects, those beliefs are the worse morally speaking, *vis-à-vis* their subjects, to the extent that there is a risk that they will be acted upon and that, when acted upon, they are likely further to entrench existing injustices. The fact that Albert does not manifest or act on his beliefs does not undermine the claim that he wrongs the subjects of those beliefs; the fact that, were he so to act, he would harm them to a greater or lesser extent lends support to the view that some doxastic wrongs are worse than others.¹⁵

IX

Conclusion. We wrong one another, doxastically speaking, when we hold about one another beliefs which manifest our failure to treat one another as autonomous and moral agents - even when those beliefs track non-spurious generalisations. The epistemic status and valence of our beliefs are irrelevant to the truth of the doxastic wrongs thesis, though they are relevant to the degree to which we wrong one another.

To some, perhaps many, the claim that the Kantian View extends to what I have called non-classic cases might seem a *reductio*. The onus is on them, thus, to show either that particularised verdictive beliefs based on non-spurious race and gender based (or similar) generalisations do not wrong the subjects of those beliefs, or to provide a different argument for the view that they do which does not extend to the non-classic cases.¹⁶

*All Souls College
Oxford, OX1 4AL*

References

- Basu, Rima. 2018. 'Can Beliefs Wrong?' *Philosophical Topics* 46 (1): 1–17.
- . 2019a. 'The wrongs of racist beliefs'. *Philosophical Studies* 176 (9): 2497–2515.
- . 2019b. 'What we epistemically owe to each other'. *Philosophical Studies* 176 (4): 915–931.

¹⁵ My argument does not have the resources to show that a racist hermit who will live in complete isolation until he dies does a greater wrong to, e.g., James in *Knife Crime* than to Warren in *Beaten Up*. My intuition is that he wrongs them to the same degree.

¹⁶ I am grateful to Kimberley Brownlee, Ruth Chang, Alexander Greenberg, Kate Greasley, Tarunabh Khaitan, Linda Eggert, Gopal Sreenivasan, and the audiences at the St Andrews CEPPA 2021 Graduate Conference, the Oxford PLP Colloquium and the Philosophy Club of the University of Washington in St Louis, for illuminating comments on earlier drafts. The penultimate draft was read at a meeting of the Aristotelian Society on November 15 2021. The ensuing discussion was enormously helpful, as were Guy Longworth's probing written comments

- Basu, Rima, and Mark Schroeder. 2018. 'Doxastic Wronging'. In *Pragmatic encroachment in epistemology*, edited by Brian Kim and Matthew McGrath. London: Routledge.
- Begby, Endre. 2018. 'Doxastic Morality: A Moderately Skeptical Perspective'. *Philosophical Topics* 46 (1): 155–172.
- . 2021. *Prejudice - A Study in Non-Ideal Psychology*. Oxford: Oxford University Press.
- Blum, Lawrence. 2004. 'Stereotypes And Stereotyping: A Moral Analysis'. *Philosophical Papers: Immoral Beliefs* 33 (3): 251–289.
- Bolinger, Renée Jorgensen. 2020. 'The rational impermissibility of accepting (some) racial generalizations'. *Synthese* 197 (6): 2415–2431.
- Brownstein, Michael, and Jennifer Saul, eds. 2016a. *Implicit Bias & Philosophy (vol.1) - Metaphysics and Epistemology*. Oxford: Oxford University Press.
- , eds. 2016b. *Implicit Bias & Philosophy (vol.2) - Moral Responsibility, Structural Injustice, and Ethics*. Oxford: Oxford University Press.
- Buchak, Lara. 2014. 'Belief, credence, and norms'. *Philosophical Studies* 169 (2): 285–311.
- Calhoun, Cheshire. 1989. 'Responsibility and Reproach'. *Ethics* 99 (2): 389–406.
- Capes, Justin A. 2019. 'Strict Moral Liability'. 36 (1): 52–71.
- Clifford, William Kingdon. 1877. 'The Ethics of Belief'. *Contemporary Review* 29:289–309.
- Eidelson, Benjamin. 2013. 'Treating People as Individuals'. In *Philosophical Foundations of Discrimination Law*. Oxford University Press.
- . 2015. *Discrimination and Disrespect*. Oxford: Oxford University Press.
- Fritz, James. 2017. 'Pragmatic Encroachment and Moral Encroachment'. *Pacific Philosophical Quarterly* 98 (1): 643–661.
- . 2020. 'Moral encroachment and reasons of the wrong kind'. *Philosophical Studies* 177:3051–3070.
- Gendler, Tamar. 2011. 'On the epistemic costs of implicit bias'. *Philosophical Studies* 156 (1): 33–63.
- Greenberg, Alexander. 2020. 'Epistemic Responsibility and Criminal Negligence'. *Criminal Law and Philosophy* 14 (1): 91–111.
- Hieronymi, Pamela. 2006. 'Controlling attitudes'. *Pacific Philosophical Quarterly* 87 (1): 45–74.
- . 2008. 'Responsibility for believing'. *Synthese* 161 (3): 357–373.

- Holroyd, Jules. 2012. 'Responsibility for Implicit Bias'. *Journal of Social Philosophy* 43 (3): 274–306.
- Khaitan, Tarunabh, and Sandy Steel. 2018. 'Wrongs, Group Disadvantage and the Legitimacy of Indirect Discrimination Law'. In *Foundations of Indirect Discrimination Law*, edited by Hugh Collins and Tarunabh Khaitan. Oxford: Hart Publishing.
- Langton, Rae. 1992. 'Duty and Desolation'. *Philosophy* 67 (262): 481–505.
- Lippert-Rasmussen, Kasper. 2011. 'We are all Different': Statistical Discrimination and the Right to be Treated as an Individual'. *The Journal of Ethics* 15 (1): 47–59.
- . 2013. *Born Free and Equal? A Philosophical Inquiry into the Nature of Discrimination*. Oxford: Oxford University Press.
- Marusic, Berislav, and Stephen White. 2018. 'How Can Beliefs Wrong? A Strawsonian Epistemology'. *Philosophical Topics* 46 (1): 97–114.
- Mogensen, Andreas. 2019. 'Racial Profiling And Cumulative Injustice'. *Philosophy and Phenomenological Research* 98 (2): 452–477.
- Moss, Sarah. 2018a. 'Moral Encroachment'. *Proceedings of the Aristotelian Society* 118 (2): 177.
- . 2018b. *Probabilistic Knowledge*. Oxford: Oxford University Press.
- Munton, Jessie. 2021. 'Prejudice as the misattribution of salience'. *Analytic Philosophy*.
- Oberdiek, John. 2017. *Imposing Risk - A Normative Framework*. Oxford University Press.
- Pace, Michael. 2011. 'The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will To Believe''. *Noûs* 45 (2): 239–268.
- Rawls, John. 1993. *Political Liberalism*. New York: Columbia University Press.
- Saul, Jennifer. 2013. 'Implicit Bias, Stereotype Threats and Women in Philosophy'. In *Women in Philosophy: What Needs to Change?*, edited by Katrina Hutchison and Fiona Jenkins. Oxford: Oxford University Press.
- Schauer, Frederick. 2006. *Profiles, Probabilities and Stereotypes*. Cambridge: Harvard University Press.
- Schroeder, Mark. 2018. 'When Beliefs Wrong'. *Philosophical Topics* 46 (1): 115–128.
- Smith, Angela M. 2005. 'Responsibility for Attitudes: Activity and Passivity in Mental Life'. *Ethics* 115:236–271.
- Srinivasan, Amia. 2020. 'Radical externalism'. *Philosophical Review* 129:395–431.
- Strawson, Peter F. 1962. 'Freedom and Resentment'. *Proceedings of the British Academy* 48:1–25.

- Washington, Natalia, and Daniel Kelly. 2016. 'Who is Responsible for This? Moral Responsibility, Externalism, and Knowledge about Implicit Bias'. In Brownstein and Saul 2016b.
- Wasserman, David T. 1991. 'The morality of statistical proof and the risk of mistaken liability. (Decision and Inference in Litigation)'. *Cardozo Law Review* 13 (2/3): 935–976.
- Zheng, Robin. 2016. 'Attributability, Accountability, and Implicit Bias'. In Brownstein and Saul 2016b.