

Discrimination: the Good, the Bad, and the Wrongful[†]

JOHN GARDNER^{*}

Abstract: This paper advances a view of what discrimination is, one emphasizing the reasons of the discriminator. It advances that view by explaining how discrimination could be objectionable when so understood. Sometimes discrimination is objectionable because it subjects a person to a disadvantage for no valid reason. However, this is not the only case of objectionable discrimination. Sometimes one should not subject someone to a disadvantage even for a valid reason. That reason is excluded from consideration. The paper outlines a view according to which the reasons that are to be excluded from consideration, under the heading of discrimination, are up for social and legal choice. Wrongful discrimination, on this view, is normally not a *malum in se* but a *malum prohibitum*.

[†] An extensively reworked version of lecture 2 of my *Diversa Lectures* entitled *Discrimination, Disadvantage, Diversity*, delivered in São Paulo in October 2016. For questions, comments, and suggestions, many thanks to Chris Essert, Ulrike Heuer, Claire Hogg, George Letsas, Guy Longworth, Tim Macklem, Conrado Hübner Mendes, Sophia Moreau, Michael Newman, Mihaela Popa, Denise Réaume, Massimo Renzo, Andrew Summers, and Lulu Weis, as well as many others at the São Paulo and London meetings. I am only sorry that many of the best questions could not be tackled within the confines of this paper, and will have to await further work.

^{*} All Souls College, Oxford, OX1 4BH; email john.gardner@law.ox.ac.uk

What is discrimination? It is possible to be discriminating in a good way. Those who are discriminating (also called ‘discerning’) are particularly sensitive to the properties that make for a better friend, or a better film, and they select accordingly. The approving use of the word ‘discriminate’ to describe such people has not died out. Today, however, the word ‘discriminate’ and its cognates are much more often used disapprovingly. ‘That’s discriminatory’ is a common way to condemn the treatment of some person or people. Some writers on the subject diagnose a conceptual shift. They say that what we mainly talk about now is discrimination in a new sense. Like exploitation or manipulation, they say, discrimination in the relevant sense is analytically bad (e.g. Halldenius 2005, p. 456). Or at any rate, its badness is analytically on the cards (e.g. Lippert-Rasmussen 2013, p. 29).

But why analytically? An alternative view is that the negative valence of the word ‘discrimination’ belongs to the pragmatics of discrimination-talk, not the semantics. Compare ‘democratic’, a word that today is widely used to convey approval of a political system. If I react to such use by saying that democracy has no redeeming features, I may be a poor judge of political systems, but it does not follow that I have a poor grasp of what other people are talking about when they talk about democracy. They can all use the word approvingly without any implication that it cannot also be used disapprovingly. Likewise, if I ask ‘what’s so bad about discrimination?’, intending to convey the contrarian view that nothing is ever bad about discrimination, I may be a moral imbecile, but it doesn’t follow that I am changing the subject. The fact that my interlocutor puts a negative spin on the word, one that I reject, is no reason to think that we are at cross-purposes about what exactly discrimination is.

Be that as it may, in investigating below what is bad about discrimination (when it is bad), and what is wrongful about discrimination (when it is wrongful), I will be putting my faith in the following ‘valence-free’ analysis of discrimination. I will call it the ‘*Discriminator’s Reasons Analysis*’ or ‘*DRA*’ for short:

To discriminate is to treat someone or something that (supposedly) has property p differently, doing so for the (supposed) reason that he or she or it (supposedly) has property p .

I say that I will be ‘putting my faith’ in this analysis because space constraints prevent me from mounting a full defence of it here. However, I do think it is a sound analysis. I hope that what I say about the badness of discrimination (when it is bad), and the wrongfulness of discrimination (when it is wrongful), will bear that claim out. To get to that stage, let me begin by clarifying a few points and allaying a few doubts about the *DRA*.

A lot of difficulties are concentrated in the word ‘differently’. Differently from what? A tempting answer is ‘differently from another, being one who lacks property p ’ (Lippert-Rasmussen 2013, 16). On this reading, discrimination has the following essentially comparative aspect. The discriminator treats a p candidate differently from a non- p candidate. The discrimination-detective in turn compares the treatment of the non- p candidate with the treatment of the p -candidate, finding that the non- p candidate (known to lawyers as the ‘comparator’) was favoured. Here you can see the beginnings of the line of thought by which laws against discrimination have come to be known as ‘equality’ laws. There are many objections to this line of thought.¹ But the first is this. For there to be discrimination, there does not need to be a comparator. The law in most jurisdictions

¹ For an excellent set of objections, see Holmes 2005.

correctly registers this, even if theoretical writing about the law often glosses over it. Here is how it is registered in the UK's most wide-ranging anti-discrimination statute:

A person (A) discriminates against another (B) if, because of a protected characteristic, A treats B less favourably than A treats or would treat others. (Equality Act 2010, s 13.)

Notice that there is no mention here of an extant person C who enjoys more favourable treatment than B. That is because in law it is enough for there to be discrimination that A would treat others more favourably if there were others, that is, if there were any candidates for treatment other than B. For B to be discriminated against there need not actually be any such candidates.

This reveals the conceptual space for what might be called 'Robinson Crusoe discrimination'. Suppose that Crusoe encounters Friday, the only other inhabitant of the island, and (departing from Defoe's story) refuses to collaborate with him because he is black. Is that discrimination? On my analysis and on the law's analysis, yes it is. There is no need for Friday to point to a comparator, a non-black person with whom Crusoe is prepared to collaborate instead. For there can be discrimination against Friday without there being discrimination in favour of anyone else. There could also be, conversely, discrimination in favour of Friday without there being discrimination against anyone else. For there need not be anyone else. That only goes to show that not all discrimination against, nor all discrimination in favour of, also counts as discrimination between. Discrimination, in short, is not essentially comparative, at least in the advertised way.²

² Macklem 2003 mounts a more radical challenge to the 'essentially comparative' view.

This fact has important consequences for how we think about discrimination, and for how we judge it. For immediate purposes, the most important consequence is this. The heavy work in differentiating discrimination from other kinds of treatment shifts from the first (main) clause of my formulation of the *DRA* to the second (participle) clause. How so? Think about Crusoe and Friday again. To know whether Crusoe treats Friday ‘differently’ when he refuses to collaborate with him, we need to ask how Crusoe would treat other imaginary candidates for collaboration, imaginary candidates who are not black but are otherwise (so far as possible) just like Friday. Why do we care to construct these imaginary candidates? Because if Crusoe would treat these imaginary non-black candidates differently, we deduce, Crusoe is refusing to work with Friday *for the reason* that Friday is black. Actually, it is misleading to call that a deduction. It is more like a restatement. To ask ‘would Crusoe have worked with Friday if only he were white?’ is (near enough) to ask ‘was Friday’s being black the reason why Crusoe wouldn’t work with Friday?’³ Asking the first question is a thought-experimental heuristic for answering the second (Holmes 2005, 186). And the second question is the one we care about when we investigate Crusoe as an alleged discriminator. To generalize, we care about *why* the supposed discriminator treated someone as he did, what properties (the law calls them ‘characteristics’) of that someone figured in the supposed discriminator’s reasoning in dishing out the treatment. Beyond that, it matters not what the treatment was. It could be anything. It need not count as any kind of mistreatment apart from the discriminator’s reasoning. It could even be fabulous treatment in all other respects. Yet it

³ Near enough? For one discrepancy, see *James v Eastleigh Borough Council* [1990] 2 AC 751. For an explanation of the discrepancy, see Gardner 1998.

becomes mistreatment under the heading of discrimination if it is treatment that should not have been dished out to the recipient for the reason that the recipient had a certain property, for instance that of being Irish or pregnant or vegetarian or an ex-offender or golf-loving or friendly to gypsies or married to a vicar.

The discrimination, we might sum up, is in the ‘why’ rather than the ‘how’ of the treatment. Or perhaps better: The ‘how’ of discrimination is in the ‘why’ of the treatment. Q: How did he treat you? A: Discriminatorily. Q: What made it discriminatory treatment to give you an extra day off? A: The fact that he gave it to me because of my [insert some property *p* here].

You can see better now why I called my analysis the Discriminator’s Reasons Analysis. Since it is (near enough) the analysis of discrimination used in anti-discrimination law, you may be surprised to learn that the DRA is unfashionable among theorists of anti-discrimination law. They tend to regard it as at odds with certain major expansions of anti-discrimination law that found juridical and political favour in the late twentieth century. They tend to think that adherents of the DRA must be (reactionary?) opponents of these expansions (Khaitan 2015, pp. 150, 152, 161). The most important such expansion is the legal recognition of the existence of ‘indirect discrimination.’⁴ Since I am by and large an enthusiast for this expansion, it falls to me to explain how it is consistent with the DRA.

The fear that the DRA tells against this expansion is not wholly baseless. Jeremy Waldron once influentially charged that casting indirect discrimination as a kind of discrimination was an instance of ‘persuasive definition’:

⁴ Beginning with Griggs v Duke Power Co. 401 US 424 (1971).

[B]y altering the descriptive meaning of a word [in this case, the word 'discrimination'], one seeks to transfer the force of its emotive meaning to a new range of actions or situations. (Waldron 1985, p. 93.)

What alteration of the 'descriptive meaning' of 'discrimination' did Waldron have in mind when he wrote this? It was the law's designation as discriminators (under the heading of 'indirect discrimination') of 'a lot of people who have no discriminatory intention' (Waldron 1985, p. 96). And what is thought to be the connection between this and the DRA? Subscribing to the DRA is thought to commit one to the idea that discrimination requires a 'discriminatory intention' (Khaitan 2015, p. 150). So the DRA must, it is thought, cast so-called 'indirect discrimination' as not really discrimination at all.

There are quite a few misconceptions here. I will mention two. The first is a misconception about where the innovation in the legal invention of indirect discrimination lies. *Pace* Waldron, it is not a semantic innovation. There is no departure from the 'old [i.e. previously prevailing] meaning of the word' (Waldron 1985, p. 93). Take this definition of indirect discrimination:

A person discriminates ['indirectly'] against a woman ... if ... he applies to her a requirement or condition which he applies or would apply equally to a man but (i) which is such that the proportion of women who can comply with it is considerably smaller than the proportion of men who can comply with it, and (ii) which he cannot show to be justifiable irrespective of the sex of the person to whom it is applied, and (iii) which is to her detriment because she cannot comply with it. (Sex Discrimination Act 1975, s 1(1)(b).)

The cumulative criteria (i) to (iii) are complex. But one need not even notice them to see that what is defined here is a kind

of discrimination. It is enough to see that the indirect discriminator ‘applies ... a requirement or condition’.⁵ In other words, he looks for a certain property *p* in his candidates – say, the property of being over six feet tall – and he treats candidates differently if and because he doesn’t find that property in them. That they do not meet his height requirement is the discriminator’s reason for treating them as he does. This is straightforward (‘direct’) height discrimination. What does it have to do with sex discrimination? This is where criterion (i) comes in. If height discrimination creates a statistically greater obstacle to women candidates, and the other criteria (ii) and (iii) are satisfied, height discrimination also qualifies as ‘indirect sex discrimination’.

It might reasonably be objected that the expression ‘indirect sex discrimination’ is too elliptical. Maybe it would not be one’s first guess that it designates ordinary (‘direct’) discrimination against people with some *other* property (people of less height, or with less availability for evening work, or having less upper body strength, or with a record of lower earnings), where that other property is statistically correlated with sex. But the problem here is only the shorthand placing of the word ‘sex’ before the word ‘discrimination’. It is not the word ‘discrimination’ itself, which is being used in an ordinary (*DRA*-conforming) way (Holmes 2007, p. 184). To simplify: it turns out that in spite of its title the Sex Discrimination Act does not regulate sex discrimination alone. It also regulates, in a derivative and circumscribed way, height discrimination, strength discrimination, and so on.

Here is a second misconception, or source of misconception. It is misleading to say that the discriminator,

⁵ Or in other legal definitions, a ‘policy’, ‘rule’, ‘provision’, ‘criterion’, ‘practice’, or ‘preference’. What I say here applies to all of these alternatives alike, in spite of other important differences between them.

under the DRA, must have a ‘discriminatory intention’. True, discriminating is something one does only in doing something intentional. Suppose that I accidentally drive into your car. That I did it accidentally, hence unintentionally, means that the possibility of my having done it discriminatorily – say for the reason that you are curly-haired, or for the reason that you are Norwegian – does not arise. That is because I did not do it for any reason. True, I did it in the course of doing something else (driving) which I was doing for a reason. Only what we intentionally do is done for reasons,⁶ and so only what we intentionally do can be discriminatorily done. Should we express this by saying that the intention of the discriminator must itself be discriminatory? No. That suggests a stronger thesis. One hears it to mean that the discriminator must intend to treat someone differently, maybe even that he must intend to discriminate. Among the earliest achievements of the courts administering modern anti-discrimination law in the UK was to establish that no such intention is needed.⁷ In law, discrimination is a strict liability wrong. That I decline to talk to you because you are a Sikh is discrimination ‘on religious grounds’ (as the law puts it) irrespective of whether I intended to treat you differently, and irrespective of whether I intended to discriminate. I need not even have the concept of discrimination (under that name or any other) to qualify as a discriminator. The discriminatory aspect of my action can in that sense be accidental even though the action in which I discriminate, whatever it may be, must be intentional. The

⁶ Note that this claim extends only to what we do (which includes, here and throughout, what we omit to do). It does not extend to what we think, want, believe or feel (in respect of which we usually have reasons but no intentions).

⁷ See R v Birmingham City Council ex parte Equal Opportunities Commission [1989] 2 WLR 520 per Lord Goff at 526: the point is ‘well established in a long line of authority’ beginning in 1980, to some of which he refers.

DRA is distorted by being recast as the view that discrimination itself can only be intentional.

II

Rudimentarily bad discrimination. To get us thinking about what makes discrimination bad, when it is bad, let me begin by saying something about justification.⁸ Here are three key points:

- (a) If there are any reasons not to ϕ , then ϕ ing calls for justification.
- (b) ϕ ing is justifiable if there are reasons in favour of ϕ ing, and the reasons against ϕ ing are not sufficient to outweigh or otherwise defeat all the reasons in favour of ϕ ing. The agent must be left with at least one undefeated reason to ϕ , if ϕ ing is to be justifiable.
- (c) ϕ ing is actually justified only if the agent ϕ s for one or more of the undefeated reasons to ϕ , and not for any of the defeated reasons.

Let me offer some commentary that will bring out the relevance of these points to our topic, beginning with a comment or two on proposition (a). The condition in proposition (a) is arguably satisfied in respect of every action. The fact that doing something will take time or effort is a reason not to do it, and every action (you may say) takes some time and effort. So proposition (a) may be held to suggest that every action calls for justification. That is not a very alarming idea. But still, it needs a bit more attention. For example: if all the options for action will take similar time and effort, it's arguable that one should no longer count the time and effort as reasons against any of those actions.

⁸ Sketching a view defended in Gardner 2007, ch. 5.

Be that as it may, in thinking about the badness of discrimination (when it is bad), we are mainly interested in a more discrimination-specific reason for the discriminator not to do what he does. We are mainly interested in the fact that what the discriminator (D) disadvantages the person discriminated against (E). D rejects E for a place on the team or a part in the show; D fails to consider E for a pay rise or a party invitation; D obstructs E in obtaining parole or winning at poker; and so on. Disadvantaging is not an essential element of discriminating, but it is an essential element of discriminating against. Discriminating against is discriminating to someone's disadvantage.⁹ We already saw, in the statutes quoted above, that UK law regulates discrimination only when it is discrimination against. As well as a D, there must be an E who suffers a detriment. The first step towards an explanation for this focus is that discrimination against is the kind of discrimination that calls for justification. There is already a distinct reason for D not to do it, and that reason is the fact of E's disadvantage that is built into the 'against'.

Now for a couple of comments on proposition (b). In formulating this proposition I was assuming that reasons often conflict among themselves. Most if not all actions have their pros and cons. In order to have a justification for doing something, one does not need to show that the pros win. That is asking too much. It is enough that the pros do not lose, in other words that the cons do not win. This assumes that at least sometimes one could reasonably have done what one did or one could have done something else instead. Reasons at least sometimes do not settle conclusively what one is to do.

⁹ It need not be to anyone's overall disadvantage. A job applicant could be lucky to miss out on a job in a failing business. Still he missed out on a job.

Different reasonable people would perhaps have done different things.

In proposition (b) I am also allowing, although I am not insisting, that there is more than one way for a reason to be defeated by opposing reasons. The utilitarians taught us to think of it all as a matter of weighing and outweighing, but the force of reasons varies in other ways too, so that they do not only compete on weight. I chose the language of ‘defeated’ and ‘undefeated’ reasons to keep my options open here. We will come back to the other option, the alternative to outweighing, in section III. But for now, I just point out that a Benthamite picture of how reasons compete is not being presupposed.

Proposition (c) is perhaps the one that requires the most explanation. To understand its appeal, let’s leave discrimination aside for a moment. Consider someone (S) who is in a situation calling for self-defence. Suppose S is being attacked by an aggrieved former spouse (T) who is intent on killing S. Using lethal force in the face of T’s attack is, let’s assume, a *justifiable* response by S. But as it happens S doesn’t realise that she is under attack. T is such a smooth operator that he makes it seem like his approach is nonthreatening, which is a stratagem to lull S into a false sense of security before plunging the knife in. Unfortunately for T, however, he misjudged what S is capable of. Unbeknown to T, S is just as embittered and filled with hatred as is T. S has been waiting for an opportunity to get revenge on T and takes the opportunity of his apparently nonthreatening approach to do so. Is S *justified* in plunging the knife into T? No. Or at any rate, the fact that she is under attack at the time doesn’t help to justify her action. That is because she did not act *in* self-defence – that is, for self-defensive reasons. For S to be justified in what she did *qua* self-defence, the fact that she was being attacked (which we are taking to be an undefeated

reason to resort to lethal force) must also have been S's reason for resorting to lethal force.¹⁰

Or take this example. An arms dealer (X) says that his death-dealing trade does not keep him awake at night because, if he didn't sell arms to the Baddies, someone else would sell more and worse arms to them. This is sometimes known as 'the arms dealer defence' and we are rightly unimpressed by it (Gardner 2007, ch 3). Why? Because it shows, at most, that X's actions were justifiable. It does not show them to be justified. We need to ask X whether he sold the arms to the Baddies in order to stop other arms dealers from doing worse. X's action moves from justifiable to justified only if preventing something worse being done was X's reason for doing what he did.

This reveals why, in discrimination situations, we are interested in D's reasons for acting to the disadvantage of E. We want to know whether D was justified in doing so. True, it would be good to reduce disadvantaging actions across the board, if only we could. But if somebody has to be disadvantaged, say because there are not enough houses or jobs or school places to go around, it had better be by way of a justified action. Disadvantaging is bad, but unjustified disadvantaging is worse. To avoid unjustified disadvantaging, D needs to disadvantage E for an undefeated reason. In the cases in which discrimination is bad, or rather in the cases in which (in the terminology I will favour) it is rudimentarily bad, that condition is not met.

How so? Because in the cases in which discrimination is rudimentarily bad, what D takes to be a reason to disadvantage P is actually no reason at all. D is mistaken to regard it as any kind of reason, never mind an undefeated one. His action for that reason is therefore unjustified even if it was

¹⁰ The famous legal illustration of the point is R v Dadson (1850) 169 ER 407.

justifiable, even if there was some undefeated reason for which he could have acted in the very same disadvantaging way. Maybe the mobility-impaired applicant was also less qualified for the auditor post, and maybe that would have been an undefeated reason for D to deny her the job. Yes, maybe. But that wasn't why D denied her the job. D denied her the job because she was mobility-impaired. D took her mobility-impairment to be a reason for not hiring her when actually it was no reason for not hiring her.¹¹

This explains why the word 'supposed' appeared in brackets before the word 'reason' in my formulation of the DRA. Recall:

To discriminate is to treat someone or something that (supposedly) has property *p* differently, doing so for the (supposed) reason that he or she or it (supposedly) has property *p*.

That people have property *p* is, in reality, no reason at all to treat them as the discriminator does. The discriminator mistakenly sees it as a reason and acts on it as if it were a reason; we call it 'her reason' with implicit scare-quotes around the word 'reason'. It is her so-called reason, her would-be reason, her supposed reason. But it is not, as we might sometimes say contrastively, a valid reason, a true reason, a real reason.

No doubt there are cases of rudimentarily bad discrimination that are just as crude as that. Q: Why did you refuse to hire this candidate? A: Because she's a woman. Q: Yes, but what's your reason for not hiring women? A: The fact that they are women is reason enough. That's a puzzling final

¹¹ This is roughly what Lippert-Rasmussen 2013, pp. 22-4, calls the 'irrelevance' account. He offers reasons to reject it as an account of what is bad about bad discrimination in general. But he offers no reason to reject it as an account of what is bad about rudimentarily bad discrimination.

answer, but imaginable. Much more routine, however, are the slightly more complex cases that are marked by the parenthetical ‘supposedly’ before ‘has property p ’ in my formulation. In these cases, the discriminator may be correct to regard having property p as a reason for the treatment that he metes out, but is mistaken in regarding the person discriminated against as a possessor or bearer of that property p . Q: Why did you refuse to hire her? A: Because she’s more interested in her family than in the job. Q: What makes you think so? A: She’s a woman, and women are more interested in their families than in their jobs.

Here the discriminator has, let’s concede for the sake of argument, a real reason, not just a supposed one, for refusing to hire some possible candidates. The property of being more interested in one’s family than in one’s job, let’s concede, could be a reason not to hire someone for a job. The discriminator’s mistake is to relegate all women candidates, and by that route this woman candidate, to the ‘more interested in family than job’ class. It is tempting to think we are dealing with indirect discrimination here. It is discrimination on grounds of interest in the job, you might say, not on grounds of sex, and it has a disproportionate impact on women candidates because of an epistemic error by the employer. It is sex discrimination in the elliptical sense I identified. But that is the wrong model. In a case of indirect discrimination against women, the employer doesn’t count the candidate’s being a woman against her, but only counts against her a different property p that (whether the employer knows it or not) happens to be more widespread among women, or among women candidates, than among their male counterparts. In the case we are now imagining, by contrast, the employer does count the candidate’s being a woman against her, and he does so in counting against her a property p that he supposes (let’s assume falsely) to be more widespread (either universal or defeasibly universal) among women, or

among women candidates, than among their male counterparts. We should regard that as a straightforward case, perhaps even the typical case, of rudimentarily bad (direct) sex discrimination. Sure, the employer may have a valid reason to treat some people differently (that they are more interested in their families than in their jobs) but he has no reason to treat this particular candidate differently under that heading. Only a false supposition about women leads the employer to hold it against her as a woman. Just as in the cruder case in which the discriminator (call him D1) holds being a woman against a woman applicant without inferring any further property p , this discriminator (D2) does not have the reason that he supposes himself to have. *A fortiori* he does not have the undefeated reason that he would need to have if he were to be justified in treating this candidate, or other women candidates, in this disadvantageous way.

If D2's is a typical case of rudimentarily bad discrimination, it is easy to see why so many are drawn to associate bad discrimination closely with prejudice (Waldron 1985, p.93; with more subtlety, Alexander 1992). D2 is prejudiced; he attributes properties to each woman candidate, because she is a woman, with inadequate regard to her actual properties (other than that of being a woman). More particularly, he proceeds by stereotyping; he attributes to each woman candidate a *standardized* set of properties that he (conclusively or defeasibly) attributes to all women, because they are women. Thinking about the hold that stereotypes may have over a person's worldview can also help us to make more sense of the strange final answer – 'the fact that they are women is reason enough' – in the case of D1. For D1, perhaps, the classification 'woman' comes, as it were, pre-stereotyped. He need not infer from the fact that a candidate is a woman that she will be more interested in her family than in her job because, for him, that is part of the very idea of a woman. It is for this notably antediluvian employer a conceptual truth that

women are more family-oriented than job-oriented. Once we understand this, we can hear 'she's a woman' as an intelligible final answer to the question of why D1 did not hire her: the fact that she's a woman entails that she is more interested in her family than in her job, and that much (we are conceding arguendo) is a possible reason to reject her application. The difference between D1 and D2, then, may lie in the level of thought at which the stereotype operates. For D1, it operates conceptually; for D2, it is a contingent (if reliable) connection, and calls for an inference. Either way, however, the discriminator exhibits his prejudice.

Yet rudimentarily bad discrimination is not the same as prejudice, and it neither entails prejudice nor is entailed by prejudice. Prejudice, like bias, superstition, and gullibility, is an epistemic fault. It manifests itself in belief. It need not manifest itself in action, let alone in the (mis)treatment of others. Take someone who, knowing of her prejudice against Muslims, only lives in towns where no Muslims live. She deprives herself of the opportunity to discriminate against Muslims (assuming that they suffer no detriment from her absence) but she does not, as far as we know, do anything to alleviate her prejudice. Conversely, those who discriminate against Muslims need not harbour any prejudice against them. Their discriminatory treatment may instead be explained by bias or superstition or gullibility. They may believe that allowing a Muslim to join one's family by marriage will bring the wrath of Kali down upon them, or they may have been duped by fake news into thinking that it is illegal to hire Muslims. No judgment, and hence no prejudgment, of any Muslim is needed to be a subscriber to these myths.

More than that: no epistemic fault of any kind is needed. A discriminator can equally be reasonably mistaken in the mistaken suppositions about reasons that bring her treatment of another under the heading of rudimentarily bad discrimination. Bad discrimination need not be the fault of the

discriminator. It may be excusable. The important question is: why would one need to excuse it? The answer is that, lacking the reason that one supposed oneself to have when one acted, one could not justify what one did. And why would one need to justify it? Because, to repeat, in discriminating against someone one is necessarily acting to the disadvantage of that someone. The rudimentary badness of what one did, *qua* discriminatory, lay in one's not being justified in acting (as UK law puts it) 'to [another's] detriment.' Acting to someone's detriment is already bad; not being justified in doing so adds the extra badness that interests us here. It is the rudimentary badness of discrimination, in the cases in which discrimination is rudimentarily bad.

III

From the bad to the wrong. It is common to hear it said, in contemporary politics, that the badness of discrimination lies in its irrationality. Talented and productive people are being overlooked! Profits are being hit! This way of challenging discriminatory practices leaves hostages to fortune. Yes, *rudimentarily* bad discrimination is irrational. It is action for what is erroneously taken to be a reason. But other bad discrimination need not be irrational. And it can be highly profitable. So the case against it must be different.

Let's return to the case of D2, the discriminator who gives as his reason for not hiring someone 'She's a woman, and women are more interested in their families than in their jobs.' We took the clause after the conjunction to be false, an instance of stereotyping. Only a false generalization about women leads D2, we said, to hold being a woman against *this* woman, E. But now imagine a more conscientious D2, call him D3, who uttered the same words. D3 protests that the clause after the conjunction is too vague as it stands to allow us to accuse him of an epistemic error. He didn't mean it the way

we read it. What he meant, more exactly, was this: women in the candidate pool for this job (likely to be at a relatively early stage in their careers and hence likely to be at a relatively young age) are more likely than their male counterparts to take extended time away from the workforce within three years to look after young children, something which is likely to be expensive and disruptive for their employers. Four claims of likelihood in quick succession! To the challenge that they are all just hunches, D3 may reply: 'Not at all. I investigated the statistics on the official Equality Agency website, no less. These things are *actually* likely, all four of them. So I did have a reason for what I did. The likelihood of a bad outcome (here, expense and inconvenience) is incontestably a reason to take steps to avoid it. It may be an open question how much weight I should attach to that reason when it tells in favour of disadvantaging a talented young woman by denying her a job that will make a big difference to her. Perhaps the extra expense and inconvenience should carry very little weight. Fortunately, I *gave* it very little weight. This candidate and the male candidate to whom I ultimately gave the job were indistinguishably excellent in every other way. It only took a tiny reason to tip the balance. And the greater probability that this candidate would be taking a career break within the next three years was that tiny reason. If all else is equal, a reason need not be very weighty to be the reason that makes all the difference.'

You may think, as I do, that the bad outcome of an action at t_2 (even if unforeseeable) is a reason not to perform that action at t_1 . Even if that is what you think, however, the relative probabilities at t_1 of different outcomes at t_2 still furnish *derivative* reasons for action at t_1 . If one has a reason to avoid being mugged, one has a derivative reason to reduce the risk of being mugged. Nobody would accuse me of erring in my reasoning if, to avoid being mugged, I also avoid going into

neighbourhoods with characteristics (being ill-lit, being full of nooks and crannies, having few people on the streets) that correlate to a higher likelihood of my being mugged. Likewise, nobody would accuse me of erring in my reasoning if, to avoid getting caught in the usual rush-hour traffic, I leave for work early – even if today there turned out to be very little rush-hour traffic. So how can D3 be accused of erring in his reasoning when, in order to avoid expense and inconvenience, he avoids hiring people with properties (being a woman, and of a certain age-group) that combine to increase the likelihood that a career break, with inconvenience and expense, is on the cards? Yet he will be so accused. He will be accused of objectionable sex discrimination against E, the objectionability of which, under the DRA, must lie in his reasons. Wherever else in his reasons it may lie, however, the objectionability does not lie in his lacking the reason that he took himself to have for refusing to hire E. For unlike the antediluvian D2, D3 did not lack that reason.

Here is another example raising the same puzzle, but without the complicating presence of any statistical generalization. It is a case of what could be called ‘inherited discrimination’.¹² D4 runs a gym. Most of his members don’t like foreigners. D4 has no idea why they take this position. He regards them as ignorant and base. Personally he would prefer a more diverse and open crowd of regulars. Nevertheless, if he extends membership to foreigners, the gym is doomed. His ignorant and base customers will desert the place in droves, and round here there aren’t enough sensible and decent types, never mind enough foreigners, to replace them. So to keep his business alive, D4 has to be unwelcoming to foreigners who come asking about membership. Let’s assume

¹² It is sometimes called ‘reaction qualification discrimination’, because it is discrimination against (or in favour of) E for the reason that others will react badly (or well) to some property *p* that E has. See Mason 2017.

he makes no mistake in how he counts the fact that they are foreigners. He has done his sums well. The problem is not that, without the discriminatory policies, he is going to be out of pocket. The problem is that, without the discriminatory policies, there is going to be no gym. And nobody gains from that, not even the unwelcome foreigners.¹³ In these circumstances the fact that someone is foreign really does count against them in the allocation of access to D4's establishments. It shouldn't count. His customers should be less foolish and base. But it does count. It counts because D4 inherits reasons to discriminate in his own actions from the prejudices of his customers. Like D3 he will be accused of objectionable discrimination, which must be, according to the DRA, a kind of error in his reasons. But wherever the error in his reasons may lie, it does not lie in his lacking the reasons that he takes himself to have.

These cases help to shift our attention from rudimentarily bad discrimination to wrongful discrimination. I don't mean to suggest that rudimentarily bad discrimination can't also be wrongful discrimination. It certainly can be. I also don't mean to suggest that there is nothing bad about wrongful discrimination. Of course there is. Analytically, it is bad to act wrongfully. But that is the badness of an action that comes of its wrongfulness. A trickier question is whether the wrongfulness of an action always comes of its (independent) badness. The cases of D3 and D4 help us to focus on this question, or an aspect of it. They are cases in which the discriminator really has the reason that he takes himself to have. If his discrimination is wrongful, that is not because it is (apart from its wrongfulness) unjustified, and hence not because it is bad in the way that rudimentarily bad

¹³ A strict egalitarian may say that there is an equality gain here, but would still be hard-pressed to show how the equality gain could also qualify as a gain for any of the people involved. See Temkin 2000.

discrimination is bad. It is not unjustified apart from its wrongfulness. Could it somehow be unjustified *in virtue of its* wrongfulness? That is an approximate rendition of the view that I favour. Let me try to make it less approximate.

Sometimes, discrimination is a social problem. The problem is that people with certain properties face discrimination frequently, on multiple fronts. When I say 'certain properties' I don't mean that they tend to be the same ones at all times and in all places. They could imaginably be any properties. One can imagine a society in which it is left-handed people, or orphans, or vegetarians, or people with tattoos, who are on the receiving end. For our purposes, it does not matter why these are the disparaged properties (as I will call them for short). What matters for our purposes (and what qualifies them as disparaged in my sense) is only that they are very widely counted against those who possess them. It may be hard for those who possess them to find work, housing, finance, education, transport, perhaps even to form personal relationships. At the very least it tends to be harder than for others. The accumulation of discrimination across various settings consigns those with the disparaged properties to poverty, unemployment, alienation, dependency, loneliness, ill-health, self-doubt, and various other disadvantages, often mutually compounding. It does not much matter which disadvantages we fix our attention on at this point. They may differ from time to time and from place to place and from disparaged property to disparaged property. The feature that matters here is the part played by widespread discrimination in creating and maintaining the pattern and severity of whatever disadvantages they are. One way of improving the lot of the folks thus disadvantaged is to tackle the widespread discrimination against them. But that task is not limited, alas, to tackling the rudimentarily bad discrimination against them, that is, discrimination by irrational disparagers. For in a society where many people are

irrational disparagers, other people – like the reluctant discriminator D4 – readily acquire reasons, sometimes weighty reasons, to play along with the irrationality. A bus driver in 1950s Alabama may have had weighty reasons (say, the conflict is turning violent) to dispatch a black passenger to the back of the bus even though he personally couldn't see what the white passengers at the front were getting so worked up about. A hotel manager may have weighty reasons (maybe he'll lose his job) to deny a booking for a wedding reception to a gay couple even though he can't see why the hotel owner has a problem with same-sex marriage. And an employer may have weighty reasons to avoid hiring women of child-bearing age even though he understands, rejects, and works hard to avoid, the gender-stereotyping that helps to make child-rearing a task predominantly undertaken by women.

The last example reveals that, in spite of the special complexities introduced by the role in it of statistical generalizations, the case of D3 is fundamentally another case of inherited discrimination akin to that of D4. Indeed the case of D3, plausibly developed, provides an interesting illustration of how discrimination begets discrimination, piling disadvantage upon disadvantage. Mothers of young children are often discriminated against on the ground of sex by those (including their own spouses and partners, other mothers who have given up work, etc.) who regard it as self-evidently the mother's role to stay at home with the children. They treat her going back to work as a reason to leave her out of social activities, to patronisingly bemoan her misfortune, etc. The effectiveness of this everyday discrimination in consigning mothers to the primary parenting role during the first few years of a child's life gives employers a genuine reason to avoid hiring younger women, and to invest less in their training and professional development, and to relegate them to less mission-critical roles, and so on. That in turn

gives spouses and partners and other familiars new reasons ('You are the one with the less important job!' 'You bring home less money!') to treat it as the mother's role to quit work and stay at home with the children and to treat the mother accordingly. And so the cycle continues.

Now we have a problem of a broadly collective-action type. How are we to break the cycle of discrimination, and thereby begin to ameliorate (or at least inhibit the perpetuation of) the attendant patterns of disadvantage? Prejudice, gullibility, bias and other kinds of epistemic irrationality may be hard nuts to crack; it is part of their nature that they are resistant to reasoning. But inherited discrimination of the kind that D3 and D4 are drawn into by the prejudices (and so on) of others may be less resistant to rational adjustment. *Ex hypothesi*, these two discriminators (and their ilk) are more reasonable. What they need is a reason not to act for what is admittedly a valid reason for acting (E's disparaged property) irrespective of the weight that would otherwise attach to that same reason. They need it to be the case that an action of discriminating that would have been justified on the weight of applicable reasons alone is unjustified once a particular reason (E's disparaged property) is excluded from consideration, and hence is no longer available to be acted on.

In short they need somehow to be given a duty not to count the disparaged property against E. Having a duty not to ϕ is having a reason not to ϕ that is both categorical and mandatory.¹⁴ If one nevertheless ϕ s, one is said to ϕ 'wrongfully'. Hence our present topic: wrongful discrimination. That the reason not to discriminate is categorical means that it does not come and go with the

¹⁴ Here I follow Raz 1977. Strictly speaking it is not the duty but the fact of the duty that is the reason. My formulation in the text skirts this issue.

changing personal goals of the person who has it. That it is mandatory means that it applies to the exclusion of some or all of the countervailing reasons (here, the reasons in favour of treating someone discriminatorily). Duties not to discriminate pick out some properties of E that D is not to count among his reasons for treating E as he does, however that may be, and not to count them irrespective of how much they would otherwise count for. Naturally they should also not be counted as supposed reasons (as they were, for instance, by our original prejudiced discriminator D1). But in a way that fact is less important. Merely supposed reasons should not be counted anyway. The existence of the duty, if it helps D1 at all, only helps him to avoid doing what he was not justified in doing anyway, quite apart from the duty, viz. counting E's being a woman against her in how he treats her. But for D3 and D4 the existence of the duty makes a bigger difference. It takes actions that D3 and D4 would ex hypothesi have been justified in performing and it makes them unjustified by excluding from consideration what would otherwise have been undefeated reasons for performing them. The duty not to discriminate, like any other duty, takes a reason that is undefeated by weight and renders it defeated by exclusion, never mind its weight. Being defeated, it is not available for D3 or D4 to act on. If they act on it they fall into rational error. That much they have in common with D1.

Perhaps I am making it sound as if all or most duties are duties not to discriminate. All duties exclude some reasons from consideration. It only requires the reasons in question to be the properties of someone or something and there we have it: a duty not to discriminate. But that move is too quick. Duties not to discriminate have an added feature. With duties not to discriminate, the reasons excluded from consideration, hence defeated by the force of the duty, also figure in specifying the content of the duty, in specifying what it is a duty to do. If I have a duty to keep my promise, some reasons

are no longer eligible to count among my reasons for breaking it. If I break the promise for one of those reasons, however weighty, I am not justified in breaking my promise. But I do not need to act for one of those reasons to put myself in breach of my duty. To breach my duty, it is enough that I don't do as I promised, never mind why. The 'why' only matters for justification. Duties not to discriminate are different in that acting for one of the excluded reasons is what puts one in breach of the duty. That is exactly what the duty is a duty not to do. I expressed that point in the following way in section I, in summarising the DRA: the 'how' of discrimination lies in the 'why' of the treatment.

When discrimination is wrongful (meaning, in breach of duty) it is analytically bad. But, assuming it is not rudimentarily bad, is it bad independently of its wrongfulness? In one respect it is, for ex hypothesi it contributes to the disadvantages associated with the possession of socially disparaged properties. That is the first step in the argument for classifying it as wrongful. But in another respect it is not bad apart from being wrongful, for had it not been wrongful, it could have been justified. It is the wrongfulness of the discrimination that blocks its being justified and lends it its specifically discriminatory badness.

IV

Wrongful discrimination in law and society. A couple of pages back I said that, to break the cycle of discrimination and tackle the attendant patterns of disadvantage, agents like D3 and D4 need to be given a duty not to discriminate. Why not just say that they have a duty not to discriminate? I am assuming that whether they have the duty depends on whether their having it would contribute to solving the collective action problem that I mentioned, namely, breaking the cycle of discrimination and ameliorating the attendant pattern of disadvantage. And I

am assuming that the making of such a contribution depends on the duty being recognized in the law and/or in social custom. With a social problem, people need to be drawn into a social solution.

Is that the right set of assumptions? That cannot be settled without more argument, and some empirical research. Here I am limited by space to offering just a little elaboration.

In the picture I paint (which may not be the whole picture) wrongful discrimination is a malum prohibitum rather than a malum in se. A malum in se is an action that is already wrongful (in breach of duty) without anyone's having proscribed it. Proscribing it is a further step. A malum prohibitum is an action that becomes wrongful only once proscribed, in other words only when we have some institutional or customary duty not to engage in it. This does not mean that there is no independent case for the people who have duties not to discriminate to have those very duties. It does not mean that, apart from the law or social custom, the duties or their assignments would have been indeterminate. Possibly, indeed, someone (say Parliament, or the courts) already had, in advance of their assignment, a duty to assign these duties exactly as they were in fact assigned.¹⁵ In the jargon, the problem that necessitated the assignment of the duties may only have been one of assurance, not one of co-ordination. But it seems to me that, in respect of duties not to discriminate, problems of both kinds are common. True, one needs to get types like D3 and D4 to perform their duties by assuring them that others will perform their duties too, and hence that the sacrifices involved will not be pointless, as well as giving types like D1 and D2 incentives to get with the programme, by

¹⁵ That brings them very close to being mala in se. Some people may draw the mala in se/mala prohibita distinction in a way that places them on the other side of the line. On the problem of how to draw the line, see Husak 2005.

adding legal and social sanctions for discriminators. But there is also the problem of working out exactly what the duties not to discriminate are to be, and on whom they will fall, if we are to optimise them as tools for collective effort against discrimination-led disadvantage.

Thus: Should a duty not to discriminate against people on the ground that they have a disparaged property (say, being black, being a woman) also be a duty not to discriminate against people on the ground that they lack that disparaged property (say, being white, being a man)? Which disparaged properties should get the protection of these duties, relatively few or relatively many, and should all get the same level of protection? Should the level of protection always (or for that matter ever) include duties not to discriminate indirectly, or should we sometimes (or for that matter always) draw the line at proscribing direct discrimination? On whom should the duties primarily fall, those best-placed to alleviate the disadvantages, or those who are the worst discriminators apart from the existence of the duties, or who? These are essentially questions of policy. The answer: whatever works best in the circumstances. And what works best in the circumstances may include obscuring the fact that those who breach these duties are committing only *mala prohibita*. It may be better if one can get people to think (mistakenly) that discriminatory actions, or some of them, are *mala in se*.

In that case it may not be enough, to make good progress with the collective action problem, to come up with legal norms. One may need the legal norms (or some of them) to become social norms too, such that non-performance comes to meet with social disapproval going beyond the ordinary social disapproval of illegality. How does one achieve this? Giving those discriminated against an individual right to sue their discriminator might help. It might help to obscure the fact that what we are ultimately interested in is tackling the wider social problem of disadvantage that discrimination

causes. Thinking further along those lines, perhaps protecting white people from discrimination as well as black, or men as well as women, would be a neat trick. It will come at relatively low cost to the disadvantage-ameliorating policy objective (few white people and men will need to call upon the protection, since their properties are not widely disparaged). Yet it will help to obscure the disadvantage-ameliorating policy objective and thereby draw people into thinking that there must be an objection to race and sex discrimination as such, irrespective of the consequences. And because of a widespread myth that moral objections are objections to certain actions as such – on principle, as it were – that may help to persuade people that race and sex discrimination are immoral, and create a social animus against them.

That is only one imaginable tactic for building social norms that support the norms of anti-discrimination law. I have no idea how effective it would be. All I can say is that something has certainly been effective. It may be a tribute to the policymakers behind modern anti-discrimination law, or it may just be the good luck of a supportive political climate in the 1970s, that strong social norms so quickly developed in support of the law. Maybe the policymakers were pushing at an open door with many discriminators like D3 and D4, possibly reluctant inheritors of the discrimination of others. Be that as it may, many people now treat race and sex discrimination as self-evident *mala in se* quite apart from the law. That is extremely good news for the policy aims of anti-discrimination law, so far as women and people with non-white racial backgrounds are concerned.¹⁶ Not surprisingly, many people with other disparaged properties want a share of the same action. They want their anti-discrimination laws too,

¹⁶ This claim is consistent with various radical critiques of anti-discrimination law according to which the policy aims are too narrow, or divert attention from deeper problems, or similar.

so that they can have discrimination against them assimilated by the same process to the same *malum in se* mythology. But can the same trick be pulled off over and over again? Or could there be diminishing returns, perhaps even a tipping point at which the *malum in se* myth, even for sex and race discrimination, is eroded by what might be called anti-discrimination fatigue? There are signs of that around us today. They help to support my general point. The politics of discrimination really are politics. The question is not, by and large, who has which nondiscrimination duties, but which nondiscrimination duties we should assign to whom. That does not mean that protecting those with widely disparaged properties against discrimination is not an imperative of justice. It merely means that the relevant questions of justice extend to: Who should be designated to do which protective work?¹⁷

Some resistance to this picture may come of the association of discrimination with racism, sexism, homophobia, and similar modern infamies. It is not clear what exactly these infamies are. The ‘-isms’ sound like ideologies or creeds; the ‘-phobias’ sounds like pathologies. Thus the words are often used today (sometimes in a deliberately inflationary and perhaps even intimidatory way) to impute a sinister ideological or credal flavour, or a belittling pathological diagnosis, to other people’s actions or beliefs.¹⁸ Partly because the currency of the words has been debased in this way, it is hard to get more precise about what they mean. They have acquired some of the quality of vulgar abuse, high on impact but low on meaning. Probably the clearest we can get on the

¹⁷ On norms of justice as norms for allocation, never mind how they are justified, see Gardner 2011, chs 9 and 10.

¹⁸ This is where Waldron’s points about ‘persuasive definition’ have more bite. Ironically, he himself is guilty of using the word ‘racist’ in just the inflationary way described: Waldron 1985, p. 95.

meaning front is to say is that racists, sexists, homophobes, etc, are people claimed to have certain misguided negative attitudes towards people with certain properties. It may be a contemptuous or a disapproving or a superior attitude. You can easily see here the association with discrimination. The fact that someone is black, or a woman, or gay, is held against him or her, both in the attitude of the racist, sexist or homophobe, and in the practical reasoning of the discriminator. But that formulation also shows how the two come apart. For the race (sex, sexuality) discrimination of the discriminator may but need not manifest the attitude of the racist (sexist, homophobe). We already saw that not all discrimination, not even all wrongful discrimination, is explained by the bias, prejudice, superstition or other epistemic faults of the discriminator. Still less is all of it explained by the discriminator's bad attitude. For not even all the relevant epistemic faults implicate bad attitudes. I need not look down on black people to prejudge them. And I need not be prejudging them, as we saw, to discriminate against them.

Many people today are wedded to the idea that racism, sexism and homophobia, understood as attitudes or sets of attitudes, are evils that transcend the contingencies of social circumstance. Even that is a mistake. They are evils mainly of our time and place, earning their role in moral thought mainly from the twists and turns of modern history. Yet it is also true that they are not evils only on condition that they are socially or legally recognised as such. Here they contrast with wrongful discrimination, which, it seems to me, is wrongful only on that condition. Possibly cases of racist race discrimination, or sexist sex discrimination, or homophobic discrimination against gay men and lesbians, form a special exception. Possibly they are wrongful because they manifest an (independently) demeaning attitude towards those discriminated against. But with most cases of discrimination,

it seems to me, that order (demeaning-therefore-wrongful) is reversed: the discrimination is demeaning only because it is (independently) wrongful.¹⁹ And it is wrongful, to repeat, only when it is socially or legally recognised as such.

V

Concluding words. Work on the theory of anti-discrimination law has struggled to honour two contrasting thoughts. One is that anti-discrimination law is a policy device for tackling collective disadvantages. The other is that discriminating (on certain grounds) is objectionable irrespective of the consequences, and that anti-discrimination law gives recognition to that fact. In this essay I showed more sympathy for the first thought. Yet I honoured the second by proposing that anti-discrimination law does not exist to tackle patterns of disadvantage full stop, but only patterns of disadvantage that are brought about or sustained by discrimination. So anti-discrimination law is what it sounds like: it is law trying to put a stop to discrimination against those with certain disparaged properties, where ‘discrimination’ carries the ordinary meaning captured in the Discriminator’s Reasons Analysis. Group disadvantage figures principally in the law’s further aim. That is the aim of ameliorating group disadvantage (the disadvantage afflicting those with disparaged properties) by reducing the prevalence of discrimination that contributes to it. If the disadvantage is not owed to discrimination, but has other causes, then anti-discrimination law is not designed to help. In that respect anti-discrimination law is no different from any other part of the law. Should we complain that, in spite of a new homicide

¹⁹ Contrast Hellman 2011, ch 1.

law, people are dying in larger numbers than ever before? Only if it is homicide that is killing them.

You may worry that what I have said here has remained aloof, explicitly aloof, from what some see as the most pressing question. What counts as a (relevant) group disadvantage? A large amount of effort has been expended, in the literature on anti-discrimination law, to isolate the kind(s) of disadvantage towards which the banning of certain kinds of discrimination should be oriented (poverty, social exclusion, low self-esteem, lack of opportunity...) and the right way of counting them (egalitarian, prioritarian, sufficientarian...). These are questions that bear on the aims of social policy as a whole. But we need not resolve them to grasp the specific aims of anti-discrimination law and policy. The specific aims of anti-discrimination law and policy are (i) to reduce the incidence of wrongful discrimination and (ii) thereby to ameliorate patterns of disadvantage that attend upon otherwise self-perpetuating cycles of discrimination. It is a boon of what I have said that it relates the wrongfulness of discrimination to the amelioration of such patterns of disadvantage, without being wedded to any sectarian account of what counts as a relevant disadvantage, let alone of how to go about counting it. Indeed it may help to explain the relatively frictionless social uptake of the norms anti-discrimination law that its aims are relatively ecumenical: they pass muster across a wide range of otherwise rivalrous explanations of what counts as a relevant disadvantage and how it should be counted.

BIBLIOGRAPHY

Alexander, Larry 1992: 'What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies', *University of Pennsylvania Law Review* 141, pp. 149-219

Gardner, John 1998: 'On the Ground of Her Sex(uality)', Oxford Journal of Legal Studies 18, pp. 167-187

____ 2007: Offences and Defences: Selected Essays in the Philosophy of Criminal Law. Oxford: Oxford University Press.

____ 2011: Law as a Leap of Faith: Essays on Law in General. Oxford: Oxford University Press.

Halldenius, Lena 2005: 'Dissecting "Discrimination"'. Cambridge Quarterly of Healthcare Ethics, 14, pp. 455-463

Hellman, Deborah 2011: When Is Discrimination Wrong? Cambridge, Mass.: Harvard University Press.

Holmes, Elisa 2005: 'Anti-Discrimination Rights Without Equality', Modern Law Review 68, pp. 175-194

Husak, Douglas 2005: 'Malum Prohibitum and Retributivism', in R.A. Duff and S Green (eds.), Defining Crimes: Essays on The Special Part of the Criminal Law. Oxford: Oxford University Press.

Khaitan, Tarunabh 2015: A Theory of Discrimination Law. Oxford: Oxford University Press.

Lippert-Rasmussen, Kasper 2013: Born Free and Equal? A Philosophical Inquiry into the Nature of Discrimination. Oxford: Oxford University Press.

Macklem, Timothy 2003: Beyond Comparison: Sex and Discrimination. Cambridge: Cambridge University Press.

Mason, Andrew 2017: 'Appearance, Discrimination, and Reaction Qualifications', Journal of Political Philosophy 25, pp. 48-71

Raz, Joseph 1977: 'Promises and Obligations', in J Raz and P Hacker (eds.), Law, Morality, and Society: Essays in Honour of H.L.A. Hart. Oxford: Clarendon Press.

Temkin, Larry 2000: 'Equality, Priority, and the Leveling Down Objection', in M Clayton and A Williams (eds.), The Ideal of Equality (London: 00000)

Waldron, Jeremy 1985: 'Indirect Discrimination' in S. Guest and A. Milne (eds.), Equality and Discrimination: Essays in Freedom and Justice. Stuttgart: F Steiner Verlag.