

Wrongs, wrongfulness, and culpability in indirect discrimination law

T Khaitan and S Steel

The final section of this paper is still to be completed.

In this paper we examine the nature of the duty-imposing norms in indirect discrimination (ID) law. We aim to clarify two main issues. First, we determine the extent to which these norms should be understood as imposing relational duties – duties owed to particular individuals or as imposing simple duties – duties owed to no one in particular – or, perhaps, both.¹ Let's call this the *duty inquiry*. Second, we clarify the extent to which considerations of culpability should be considered to be aspects of the wrong(s) of discrimination. Let's call this the *culpability inquiry*. This paper is motivated by a recent trend in British and American law to doubt the legitimacy of the prohibition on ID and to attempt to curtail its operation.² We show through these inquiries that the prohibition on ID is legitimate.

The first part of the paper provides a general explanation of the concepts of relational and simple duties. The second and third parts undertake the *duty inquiry* and *culpability inquiry* respectively. We find that indirect discrimination usually entails the breach of two distinct duties, making it wrong and wrongful. We also claim that while indirectly discriminatory acts can be wrongful though not otherwise culpable, culpability considerations are nonetheless relevant to supporting the liability regime for indirect discrimination. The fourth part offers observations on the structure of discrimination law in light of our analysis.³

I. Duties differentiated

Let's begin with moral duties. Some moral duties are *relational* in the sense that they are owed to others.⁴ A has a moral duty not to touch B without B's consent. A owes this duty *to B*. If A breaches a duty owed to B, then A *wrongs* B. Other moral duties are not in any way relational: they are not owed to others. C has a moral duty not to waste away C's life in perpetual intoxication. But C need not owe this duty to any other person in particular.⁵ These can be called *simple* duties. If C breaches such a duty, we will say that C behaves *wrongfully* (but does not commit a *wrong*). The same action may be mandated by both relational and simple duties. So my polluting Esther's river may be prohibited both by a relational duty, owed to Esther, not to interfere with Esther's property *and* a non-relational duty not to degrade the environment, owed to no one else in particular. Where one and the same act, or act-type, is prohibited by both a simple duty and by a relational duty, we will say that there is a *composite* duty prohibiting that act.

¹ See below, X.

² Ricci, Essop.

³ This fourth section has not yet been completed.

⁴ The terminology differs here. Some refer to 'directed' duties: e.g. S May, 'Directed Duties' (2015) 10 *Philosophy Compass* 523: 'Directed duties are duties that an agent owes to some party'. We borrow the terminology of 'relational' and 'simple' duties from J Goldberg and B Zipursky, 'Rights and Responsibility in the Law of Torts' in D Nolan and A Robertson (eds), *Rights and Private Law* (Hart, 2014).

⁵ Imagine C is an orphan so has no familial obligations, pays taxes, donates to charity etc.

In virtue of what is a duty relational? One view is that is in virtue of the duty being correlated with a claim right.⁶ The nature of relational duties then depends upon the best theory of the nature of moral claim rights. Here we suggest tentative sufficient conditions for the relationality of a moral duty, while remaining neutral on whether this relationality also amounts to the best theory of moral claim rights.

A moral duty will be owed by A to B if:

- (i) *Interest relationality (IR)*. An aspect of B's well-being (an 'interest' of B's) suffices to justify B's being under the duty⁷ or
- (ii) *Status relationality (SR)*. B's moral status suffices to justify A's being under the duty.⁸

For instance, A owes B a duty not intentionally to damage B's body. This is explained either by the fact that B's interests are sufficient to justify this duty or by B's moral status (as, for instance, an autonomous agent, or the kind of thing which is capable of autonomous action). By contrast, A's duty not to spoil the environment for distant future generations is not justified by the interests or status of B, if B is a particular individual living in the same time period as A.

IR and SR come apart in cases where it seems that A's duty is not justified by an aspect of B's well-being. Kamm writes: "Persons might have a right to treatment as equals ... without our duty to them being based on their interests. Rather, I would say, this right is based on their nature as persons and not necessarily related to any aspect of their well-being. Even if it turns out to be in their interest to have this nature, the right derives from their nature and not from their interest in having it".⁹ The wrong involved in some cases of direct discrimination, which appears to be a relational wrong, may be an instance of this. If A knowingly rejects B's job application on grounds of B's race, knowing B is indifferent between obtaining this job and another job which A knows B will obtain, the discrimination may have no effect on B's well-being and yet it still seems wrong.¹⁰

What unites IR and SR is that the relationality of the duty consists in some property of the counterparty being sufficient to justify the duty. Notice that this leaves open the possibility of composite duties. The fact that a property of an individual *suffices* to justify the duty leaves open the possibility that there are other (even sufficient) justifications for the duty.

⁶ See generally on claim rights, W Hohfeld, 'Some fundamental legal conceptions as applied in judicial reasoning' (1913) 23 *Yale Law Journal* 16.

⁷ See J Raz, 'The Nature of Rights' (1984) 93 *Mind* 194.

⁸ See FM Kamm, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (OUP, 2007) 246

⁹ FM Kamm, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (OUP, 2007) 246. See also L Wenar, 'The Nature of Claim-Rights' (2013) 123 *Ethics* 202.

¹⁰ For a case arguably of this kind, see K Lippert-Rasmussen, *Born Free and Equal: A Philosophical Inquiry into the Nature of Discrimination* (Oxford, 2014), 157.

In short, then, we will assume that moral relationality is fundamentally a matter of the kinds of available justification for the duty. If the duty is justifiable by another's interest or moral status, then the duty will be owed to that individual.

Legal duties can similarly be relational or simple in character. Private law duties are perhaps the paradigmatic example of relational legal duties. My legal duty not to touch you without your consent is a duty owed to you, breach of which amounts to a *wrong*. Legal duties not to litter or to pay taxes are, by contrast, duties not owed to particular persons and thus simple duties, breach of which amounts to wrongful conduct, without a *wrong*.

A legal duty is what, from the law's perspective, one has a moral duty to do.¹¹ Therefore, a relational legal duty is what, from the law's perspective, one has a relational moral duty to do. In so far as the law imposes a duty, how does one determine its perspective upon whether the duty is relational, simple, or composite? The language used by judges and in legislation has some force here. If the duty framed in explicitly relational terms, then we would need reason to believe that the law is using relational concepts in a specialized way which does not track their ordinary moral usage. However, the language in which a duty is articulated is not decisive. What ultimately matters is whether the law regards an individual's interest or status as a sufficient justificatory source of a duty.

Quite often it will not be readily apparent what the law's perspective is on the justificatory basis of a duty that it imposes. However, by regulating duties in different ways, the law's view may be implicitly determined. A legal duty is typically relational (and not necessarily to the exclusion of its being a simple duty too) where the following are jointly instantiated: (i) a legal power to bring proceedings for the enforcement of that duty is conferred upon an individual, (ii) the proceedings, once brought, remain within the control of that individual and (iii) the person in whom (i) is vested has an entitlement to obtain (usually compensatory or injunctive) reliefs for the (prospective) breach of the duty, where they may enjoy the benefits of those reliefs.¹² These features are readily explicable by the fact that the justificatory source of the duty lies in the individual's own interests or status.

II. Relational and simple duties in British indirect discrimination law

ID is listed as a 'Prohibited Conduct' under the UK Equality Act 2010 and is defined under s. 19(2) of that Act:

“... a provision, criterion or practice is discriminatory in relation to a relevant protected characteristic of B's if—
(a) A applies, or would apply, it to persons with whom B does not share the characteristic,

¹¹ Here we follow J Raz, 'Legal Rights' (1984) 4 *OJLS* 1.

¹² Cf HLA Hart, 'Bentham on Legal Rights' in HLA Hart, *Essays on Bentham* (Oxford, 1982) where different types of control over the enforcement of a duty are distinguished and used to analyse the concept of being a right-holder. In our view, such powers of enforcement are explained by the relationality of the obligation, rather than the explanatory relation running the other way. However, the conferral of these legal powers may be good evidence of the relationality of the duty in question.

Please do not cite or circulate without the authors' permission.

- (b) it puts, or would put, persons with whom B shares the characteristic at a particular disadvantage when compared with persons with whom B does not share it,
- (c) it puts, or would put, B at that disadvantage, and
- (d) A cannot show it to be a proportionate means of achieving a legitimate aim.”

Read alongside judicial pronouncements on ID liability, clause (a) can be taken to establish the requirement that the provision, criterion or practice (‘act’, for short) is facially neutral, i.e. it is applied to all persons regardless of their race, sex or other protected characteristic (hereinafter, ‘grounds’). Implicit in this clause is the requirement that the act must not intentionally be based on any of these grounds. Any discriminatory animus will make a case one of *direct* discrimination.¹³

In the complicated s. 19(2)(b), notice first that the word “puts” establishes a causation requirement: the act in question *puts* (or would *put*) the relevant persons at a particular disadvantage.

Next, if the relevant ground is race, “persons with whom B shares the characteristic” could be black people, white people, non-whites (s. 9(4)), South Asians, Poles (s. 9(1)(b)), or any other racial group (s. 9(2)(b) Equality Act 2010). The “disadvantage” suffered by the group must be comparative: “when compared with persons with whom B does not share [the characteristic]”. The comparison required is between two groups defined by the same protected characteristic: whites and Asians, Jews and non-Jews, men and women, wheelchair users and non-users, and so on.

Finally, the provision requires proof of “particular” disadvantage. The term is borrowed from EU law (art. 2(2)(b) Framework Directive 2000/78/EC), and was intended as a claimant-friendly change from the original language requiring *disproportionate* impact (used, for example, in s. 1(1)(b)(i) Race Relations Act 1976). As the Supreme Court has recognised:

“the new formulation was not intended to make it more difficult to establish indirect discrimination: quite the reverse... It was intended to do away with the need for statistical comparisons where no statistics might exist. It was intended to do away with the complexities involved in identifying those who could comply and those who could not and how great the disparity had to be.”¹⁴

Whether framed as “disproportionate impact” or comparative “particular disadvantage”, the requirement is essentially one of correlation. When those disadvantaged by an act *disproportionately* or *particularly* belong to a protected group (say women), clause (b) has been satisfied. It is also worth pointing out that if everyone disadvantaged by an act belong to one protected group (women) and everyone advantaged or not affected by the same act belong to another protected group defined by the same ground (here, men), British law treats the act as directly

¹³ *CHEZ Razpredelenie Bulgaria AD v Komisia za zashtita ot diskriminatsia*, case C-83/14; [2015] I.R.L.R. 746 at [75], [76], [95]. Cf *Walker v Hussain*.

¹⁴ *Homer v Chief Constable of West Yorkshire* ([2012] UKSC 15; [2012] I.C.R. 704 at [14].

discriminatory, irrespective of whether there was any intention involved.¹⁵ However, if all members of the disadvantaged group suffer, but so do some members of its comparator group, the discrimination is still *indirect*, albeit “a form of indirect discrimination which comes as close as it can to direct discrimination”.¹⁶ The law, therefore, recognizes that ID liability admits to degrees of severity—greater the correspondence between the two sets of groups (those disadvantaged and advantaged by the act and those belonging to a protected group and its corresponding comparator group), more disapproving is the law of the conduct in question.

To summarize, *prima facie* ID in British law occurs when the following two requirements are satisfied:

- a causal connection between the defendant’s facially neutral act and some adversity suffered by a set of persons
- a significant correlation between membership of this set of adversely affected persons and membership of a *group* protected by discrimination law

Our claim in this section is that *prima facie* ID, at least in archetypal cases, amounts to the breach of a simple duty and a relational duty. We do not propose to outline the precise shape of these duties. That will require a thicker normative account of discrimination law than we intend to propose. We will, instead, draw up the broad contours of these duties, and point to existing normative accounts that pay attention to them. Broadly, the simple duty is characterized as the duty to refrain from causing adversity to protected groups.¹⁷ The relational duty, we will claim, is the duty to refrain from causing disadvantage to a person because of her membership of that group.

Before we clarify these duties, let us consider an example in order to understand the relationship between correlation and causation in this area of law. We choose as our exemplar case one where there is clear evidence of a significant correlation between group membership and the disadvantage, but no direct causal explanation.¹⁸ We exclude cases where there exists a direct causal explanation—e.g. between a turban ban’s impact in Sikhs or an inflexible working hours policy’s impact on women—because they are theoretically easier to defend. This correlative relationship is key to understanding the two duties. Imagine a facially neutral university admissions test, one that is not designed or administered with any discriminatory animus. In the first scenario, no particular racial group disproportionately fails the test, that is, the chances of one’s passing the test do not correlate with her race in this case. In the second scenario, black candidates are only 25% as likely as their white counterparts to pass the test, and this difference between the two groups is statistically significant. Technically, an outcome is ‘significant’ if it is improbable that it occurred by chance.

¹⁵ *Bull v Hall* [2013] UKSC 73; [2013] 1 W.L.R. 3741 [19]: “I take there to be direct discrimination when the category of those receiving a certain advantage and the category of those suffering a correlative disadvantage coincide exactly with the respective categories of persons distinguished only by applying a prohibited classification”.

¹⁶ *Rodriguez v Minister of Housing* [2009] UKPC 52; [2010] U.K.H.R.R. 144 [19].

¹⁷ Most of the discussion in this paper is in relation to protected groups that suffer abiding, pervasive and substantial disadvantage in their societies. We take these to include women, racial minorities, gays and lesbians, transpersons, and disabled persons, among others. Whether, and when, socially dominant groups should be protected requires a separate argument: one of us has provided one such argument at Tarunabh Khaitan, *A Theory of Discrimination Law* (Oxford University Press 2015) pp 171-180.

¹⁸ This closely tracks the facts of *Essop*. See LQR case note.

Imagine also a black candidate, Ifemelu, who fails the test in both scenarios. We will look into when and why Ifemelu ought to be able to bring a claim when discussing the relational duty. For now, it suffices to notice that Ifemelu (or anyone else, for that matter) will have an ID claim only in the second scenario.

Our first scenario is what statisticians call the ‘null hypothesis’—all else being equal, this is what one would expect. In our second scenario, the null hypothesis is disproved. In other words, the finding of a statistically significant difference between the two racial groups makes it improbable that the difference is down to chance or luck. The term ‘improbable’ is used advisedly. A finding of statistical ‘significance’, absent an explanation, cannot rule out chance, it can only make it highly improbable. This should not bother lawyers, who are used to working with proof on balance of probabilities. A significant correlation therefore points to *some* non-random cause of the disparity between the groups.

It is true that the law no longer requires a strict proof of statistical significance. But legal truths cannot be (and ought not to be) as fine-grained as scientific truths. Meeting scientific standards of truth imposes significant costs, both on adjudicatory institutions and on the parties (especially on the party that bears the burden of proof). To keep these costs manageable, law has to make a judgment about the probability and moral tolerability of false positives vis-à-vis false negatives. In criminal law, for example, false positives are anathema, and rightly so. It is morally much worse to convict someone who is innocent than to acquit someone who is guilty. This is why the burden of proof is typically on the prosecution and the standard of proof is very high.¹⁹ In discrimination law, litigation takes place under conditions of significant resource and information asymmetry in favour of the defendant. The legal shift from the requirement of ‘disproportionality’ to that of ‘particular disadvantage’ is in light of the recognition of this asymmetry. The law no longer requires the claimant to prove statistical significance of the differential impact. But the claimant still needs to make a plausible *prima facie* case that the difference is unlikely to be random in order to establish *particular disadvantage*. It remains open to the defendant to rebut this *prima facie* case by showing that the disparity is indeed down to chance. The divergence from the exemplar case is down to how litigation operates in the real world. But this divergence does not take away the explanatory force of that case. Even the rule of thumb inquiry that the law now permits is still asking the same question: is the difference between the groups likely to be down to chance or not.

Once chance is ruled out and the existence of *some* non-random cause of the disparity is established as a legal fact, discrimination law is best understood as presuming two causal claims:

- (i) the defendant’s act caused a particular disadvantage to blacks as a group, and
- (ii) the claimant suffered some adversity because of her membership of her racial group.²⁰

¹⁹ Amit Pundik, *Freedom and Generalisations*.

²⁰ These presumptions are made explicit in s 136(2): “If there are facts from which the court could decide, in the absence of any other explanation, that a person (A) contravened the provision concerned, the court must hold that the contravention occurred.” On shifting of the burden of proof to the defendant, see also, *Hoogendijk v The Netherlands*, 58641/00 (2005).

So long as the claimant is a typical member of the group, the second causal claim is entailed in the first. If the first claim is established, at least some members of the group suffered adversity because of their membership of the relevant group: the only remaining question is whether the claimant is one such member. The first of these causal presumptions relates to the simple duty and the second to the relational duty. But before we make that connection clearer, the drawing of these causal presumptions from the proof of a significant correlation calls for justification. For, this move from significant correlation to a presumption of causation is not uncontroversial. Even the strongest of correlation claims, without more, can rule out chance only probabilistically. We have seen that law's move from statistical proof to the rule-of-thumb standard of 'particular disadvantage' was justifiable because of the civil nature of the liability in discrimination law, and because of the information asymmetry that usually afflicts these cases. These two factors also lend partial support to the presumption of causation from the proof of correlation. But there are other reasons too, which further support this legal presumption.

One of them is that the causal presumptions are rebuttable by the information-rich defendant.²¹ It remains open to the defendant to show that despite the extant significant correlation, there is no 'particular *disadvantage*' caused to the group. In the example we are working with, assume that the defendant in question is an affirmative action employer who especially encourages black candidates to apply for the concerned job. If this is the case, it is plausible that more under-qualified black candidates end up applying for the job than under-qualified white candidates. This is likely to result in the correlation that we noticed. If the defendant can demonstrate any such benign explanation for the correlation, there is no *disadvantage* caused to the group. Then, the presumption of 'particular disadvantage' to the group will be rebutted.

The defendant can also challenge the finding of '*particular disadvantage*' by denying its particularity, despite the existence of significant correlation. She can do so, for example, in cases where the statistical outcomes are subject to the Simpsons Paradox. This Paradox was observed in graduate admissions at the University of California, Berkeley in 1973. Although the aggregate data did show a statistically significant bias against women candidates, closer inspection of the data revealed that women were much more likely to apply to departments which were more difficult to get into (for persons of either sex).²² In such cases, the defendant can effectively show that by moving the analysis to a more specific level, the correlation demonstrating particular disadvantage disappears. There may be other similar avenues available to the defendant to challenge the prima facie case of particular disadvantage.

As soon as the defendant rebuts the first causal claim, the second—entailed—causal presumption is automatically rebutted. If the defendant fails to offer any such rebuttal, it seems fair to draw the most plausible conclusion from the correlation established by the claimant: that the defendant's act caused particular disadvantage to the group in question. This is because we already have very good reasons to rule out chance. If

²¹ S 138(4) recognizes this information asymmetry by allowing a court to draw adverse inferences from a defendant's refusal to answer a question, or from his or her providing an evasive or equivocal answer.

²² P Bickel, E Hammel and J O'Connell, "Sex Bias in Graduate Admissions: Data from Berkeley" 187 *Science* (1975) 398.

there is a benign explanation for the correlation, the defendant is best placed to offer it. If no such exculpatory explanation is forthcoming from the defendant, the weight of evidence implicates her. In such circumstances, to demand that the claimant prove not merely the correlation (and the consequent causation), but also explain the cause,²³ is asking for too much.

Furthermore, even after all this is established, the defendant still has the opportunity to *justify* her act under section 19(2)(d) by showing that it is a proportionate means of achieving a legitimate objective. Even if an employer's act is found to amount to unintentional ID, declaration and recommendation are the preferred remedies under section 124 of the Equality Act, with damages permitted only exceptionally. The recommendation is likely to ask the employer to change her discriminatory and unjustified policy so that it becomes non-discriminatory, or, if discriminatory, is a proportionate response to her legitimate objectives. The multiple legal presumptions against the defendant seem a lot less worrying when judged alongside a remarkably light-touch remedial scheme.

We can now examine how the first causal claim relates to the simple duty. What has been established in law is that the defendant's act caused a particular disadvantage to the relevant group qua group. In the second scenario, therefore, Ifemelu's group—blacks—was adversely affected *as a group*. Recall that the breach of a simple duty is *wrongful*, although it does not (necessarily) wrong any particular person. The essential feature of the simple duty breached by indirectly discriminatory conduct is that it adversely affects a group that is protected by discrimination law, qua its being a group and is wrongful for that reason. This claim assumes a direct correspondence between adverse effect on a sub-set of a group and on the group as a whole. The extent of adversity on the group because of an act of ID may be trivial, or significant. An employer with a small workforce is likely to make only a minuscule impact on women as a group through indirect sex discrimination. But an indirectly discriminatory law of general application (say, the criminalization of sodomy) could have pervasive and serious impact on the entire group (in this case, gay men).²⁴ We take the assumption that there is some knock-on effect on the group as a whole to hold in most, if not all, cases. This will suffice in order to justify the *general* claim that ID breaches the simple duty.

We have only outlined the broad contours of the simple duty. A full account will need to defend it normatively. Any such normative defense is likely to account for the importance of groups protected by discrimination law, and why we have reasons to refrain from inflicting any disadvantage upon them, qua group. It will also need to explain, among other things, the nature of the adverse effects that count, and the type of groups that are/ought to be protected.²⁵

Now we can move to the relational duty. A discriminator [X] *wrongs* a particular person subject to ID [Y] because X makes Y suffer some adversity because of her

²³ As the Court of Appeal did in *Essop*.

²⁴ *Doyle*.

²⁵ Some theorists who have provided such accounts include Collins, Sunstein, Hellman, Khaitan, Koppelman, and Siegel.

membership of a protected group.²⁶ This is exactly what happens when the second causal claim—the claimant suffered adversity *because of* her membership of her group—is satisfied. The phrase ‘because of’ one’s race etc is ubiquitous in discrimination law, and has given rise to much controversy.²⁷ Most of this controversy arises as “these phrases can refer either to a) the reasons that guide the acts of agents or to b) factors that do not guide agents but do help explain why the disadvantageous outcomes of certain acts and policies fall disproportionately on certain salient groups.”²⁸ The sense in which we are using the term is obviously the latter, and this is the sense that reflects the state of British law much better than the first sense.

The second causal claim is entailed in the first for a simple reason. If the defendant caused some adversity to a group qua group, then it follows (short of any separate reason to believe otherwise) that affected members of that group suffered because of their group membership. ID breaches a relational duty, inasmuch as it entails a *wrong* to these particular members of the group. In our first scenario, although Ifemelu herself failed the test, her group (blacks) had not suffered particular disadvantage as a group. For this reason, her failure could not be connected with her group membership. In the second scenario, however, Ifemelu fails as a member of a particularly disadvantaged group. Unless there is any reason to believe that she is an atypical member of the group, her failure can normally be attributed to her group membership.

One might object to this claim by suggesting that the facts at hand are not sufficient to get to the second presumption that the claimant suffered some adversity because of her membership of her racial group.²⁹ All we know about the scenario under consideration is that black candidates are only 25% as likely as their white counterparts to pass the test, that this difference between the two groups is statistically significant, and that Ifemelu is a black applicant who failed the test. The objection is that Ifemelu might have failed the test anyway, even in this second scenario. This would especially be the case if the pass rate (for everyone) is quite low: say, the ratio between the applicants to available places is rather large. At least in such cases, it may be argued, one cannot say that Ifemelu suffered *because of* her race.

The objection is a good one, but is based on the (mistaken) presumption that the only relevant adversity that Ifemelu suffers is actually failing the test. Let us try to get inside Ifemelu’s mind to see if she suffers in the case at hand because of her race. To avoid any suggestion that this thought experiment imports subjective feelings into indirect discrimination analysis, let us ask not what Ifemelu is *likely* to think or *actually* thinks, but what would be *reasonable* for her to think. Let us also go back in time when our reasonable Ifemelu is considering whether to apply for the job in question. As a reasonable person who does not much fancy wasting her time, she is likely to apply only if she considers her chances of getting the job to be at least

²⁶ Whether a person is allowed to bring a claim of indirect discrimination is related to, but not determined by, whether they suffered a relational wrong. There may be good institutional reasons for extending standing to those who did not suffer such wrong, or restricting it to a sub-group of those thus wronged. See CHEZ.

²⁷ Gardner, Finnis

²⁸ Altman <http://plato.stanford.edu/entries/discrimination/>

²⁹ We are grateful to David Edmonds for highlighting this objection.

plausible. If she knows that as a black person, the chances of her getting the job are 25% lower than if she were white, this is sufficient for her to reasonably think that the likelihood of her getting the job depend, to some extent, upon her race. It is true that the plausibility of her getting this job will also depend, in addition to her race, on the overall success rate. Even so, a lower overall success rate in fact makes it *more* likely that Ifemelu, as a black applicant, won't get the job. If the job had a high success rate, Ifemelu's application may be plausible despite the statistical significance of race. However, the same significance in a job with a very low success rate is likely to lead Ifemelu to conclude that even making the application would be a waste of her time. Whether, for that reason, she chooses not to make the application at all, or whether, having made the application, she fails, she would have suffered (in part) because of her race.³⁰

On this view Ifemelu suffers because of her race not in virtue, necessarily, of her not getting the job, but rather because her chances of obtaining the job have been significantly reduced. Although the absolute level of her chance of the job was already low, it is made even less likely by the policy in question. The disadvantage that Ifemelu suffers here can be conceptualized in at least two ways. Following Moreau, we might conceive of it as an interference with Ifemelu's deliberative freedom. Since Ifemelu's race affects her probability of success, it thereby becomes rationally salient in her deliberations. This analysis plausibly identifies a disadvantage in cases where the victim is aware of the effect her race has on her chances of obtaining the job. We are less persuaded of the force of the analysis where the victim is unaware. This is not because we do not believe one can suffer disadvantage without being aware of it. It is rather because the focus upon the victim's hypothetical deliberations, or the deliberative perspective of a rational or reasonable victim, in cases where the victim is not aware of the probabilities *ex ante*, seems somewhat to deflect from a more basic way in which the victim suffers disadvantage in these cases. Disadvantage is suffered in virtue of the fact that one of the victim's opportunities becomes rationally less attractive and so less valuable to her. Since autonomy requires the possession of a wide set of valuable options, the diminishment in the value of an option has an impact on an agent's autonomy. It seems clear that one's autonomy can be diminished without one's knowledge: if the conference organisers lock us in the room, without our knowledge, we have been deprived of our freedom to leave, even if this remains unknown to us.³¹

Again, we have merely sketched the broad contours of the relational duty. A full account of the relational duty will need to explain why memberships of certain groups should be irrelevant to certain types of outcomes. It should also give an account of the groups whose membership is normatively irrelevant in this sense.³² Furthermore, a successful account of discrimination law should account for both types of duties breached by ID.

III. Culpability and blameworthiness in indirect discrimination

³⁰ In fact, on this account, she suffers even if she eventually gets the job. That the law does not give her standing to bring a claim in such cases is a distinct policy choice.

³¹ There may be yet be a still more basic sense in which Ifemelu is harmed. It is at least arguable that merely being at risk of disadvantage is itself a form of disadvantage: see S Steel, *Proof of Causation in Tort Law*, ch 6.

³² Such an account is provided by Moreau, Khaitan, Rassmussen.

In UK law, the ID duty (ID duty) may, on the face of it, be breached without the agent being culpable in the sense of being blameworthy. The duty (s.19) is framed without an explicit culpability requirement.

The ID duty appears, then, to be a *strict* duty. A strict duty is a duty whose breach is not contingent upon the prohibited conduct being done with fault. It is liability *regardless* of fault.³³ That is: it is liability which attaches to an agent, A, for X-ing regardless of any steps A took or could have taken to avoid X-ing and regardless of whether A knew or had reason to know whether A would X.³⁴

One's immediate impression that the ID duty is strict in this sense needs at least one qualification, however. The ID duty is not breached *regardless* of fault. This is because the ID duty cannot be breached where A acts with the purpose of causing B to be at relative disadvantage because of B's protected characteristic. If A acts with this intention, then A breaches the *direct* discrimination duty.³⁵ If the law is that direct and ID are mutually exclusive, then it follows that ID cannot be committed with the most serious form of culpability.³⁶

It might also be questioned where s.19(2)(d) has the effect of introducing culpability considerations into the question of whether a person breaches their ID duty or into the question of whether a person can justify or excuse their breach of that duty. If A could avoid being liable by demonstrating that A 'reasonably believed' that the means adopted were proportionate to achieving a legitimate aim, then it would follow that the absence of culpability would act as a valid excusatory condition. But demonstrating a reasonable belief in the proportionality of the means is clearly insufficient for it to be the case that A can show it to be a proportionate means. Showing that it *is* proportionate is obviously different from showing that one had a *reasonable belief* that it was proportionate.³⁷

The valid point remains that the ID duty can be breached *without any* fault. In this part of the paper, we examine (a) the extent to which the law, in imposing this kind of duty, departs from any plausible conception of the moral wrong(s) involved in acts of ID and (b) whether this form of strict legal duty can (nonetheless) be justified.

It will be useful to set out here four different ways in which it might be claimed that considerations of culpability figure in the moral wrong of ID:

³³ Liability and duty

³⁴ This definition is drawn from J Gardner, 'Some rule of law anxieties about strict liability in private law'.

³⁵ While the Employment Appeal Tribunal has held in *Walker v Hussain* [1996] IRLR 11 that indirect discrimination can be intentional, the European Court of Justice has recently suggested in *CHEZ* that the presence of relevant intention will necessarily amount to direct discrimination.

³⁶ In so far as one believes that ID is a truly distinctive wrong from DD (rather than simply a less culpable form of DD) this seems to be an odd result. Are there any other wrongs which cannot be committed *intentionally*? Negligence (in law) is not an example: intentionally risking another harm is just an aggravated form of unreasonable risk imposition.

³⁷ Although some earlier decisions did appear to incorporate a reasonable belief limitation here, see: S Fredman, *Discrimination Law* 2nd edn (Oxford, 2011), 193-196.

(i) *Culpability as constituent of the wrong.* On this view, culpability is part of the primary moral ID duty. There is no wrong of ID without culpable conduct.³⁸ Compare the moral wrong of negligence. On one view of this wrong, or of one species of it, it specifies a duty to have a certain attitude of care towards the interests of others. On that understanding, the wrongness of negligence cannot be specified independently of culpability.³⁹

(ii) *Culpability as necessary for responsibility for a wrong.* On this view, culpability does not figure in the primary moral ID duty, but rather enters into a specification of the conditions under which someone is morally responsible for a breach of their ID duty. It may be, for example, that the insane are capable of breaching serious moral duties (not to kill, for example) and yet are not morally responsible for those breaches.

(iii) *Absence of culpability as an excusatory condition.* On this view, culpability does not figure in the primary moral ID duty, but rather is relevant to whether the wrong was excused. A person can be morally responsible for a wrong despite being excused for it. For instance, a person might be excused for acting in the face of extreme provocation (in response, say, to domestic violence) and yet be morally responsible for their wrongdoing.

(iv) *Culpability as an aggravating feature of the wrong.* On this view, there can be non-culpable wrongs of ID, but culpability enhances the blameworthiness of the wrong, perhaps rendering the wrongdoer liable to different normative consequences to non-culpable ID-feasors.⁴⁰

(i) Culpability as constituent of the wrong

In relation to (i), one might wonder whether there is a *general* theoretical problem with such accounts. There are well-known difficulties, familiar from the literature of the doctrine of double effect, with the idea that the intention with which a person acts can affect the *permissibility* of their action.⁴¹ How can changing one's intentions change the permissibility of one's action? To give a well-known objection: It seems odd that a person advising a doctor whether to administer morphine which will shorten the life of a patient should tell the doctor that the permissibility of her action will depend upon what her intentions will be (the intention to kill or the intention to relieve suffering). Perhaps this problem generalises with other forms of mental state which could be relevant to an agent's culpability. It would be odd, too, perhaps, if the advice depended upon what the agent's mental states or attitudes other than intentions were. Strictly, however, friends of (i) probably need not be committed to the idea that culpability affects the permissibility of an action. They might think that culpability is a necessary ingredient of the wrong without thinking that its wrongfulness is

³⁸ We could further distinguish: culpability as a constituent of the wrong of ID and culpability as a constitutive of the wrongfulness of acts which amount to ID. It could be that acts of ID are wrongful independently of culpability, but that culpability is what makes them wrongful acts of ID.

³⁹ For a different understanding of negligence, see below, X.

⁴⁰ We could also distinguish between culpability's role in enhancing *blameworthiness* and its role in altering the *wrongfulness* of the act of ID. It might be that blameworthy ID is not just, well, blameworthy ID, but a different *form* of the wrong of ID. Compare the difference between killing and murder. Murder is not just blameworthy killing, it is killing with another dimension of wrongfulness.

⁴¹ J Thomson, 'Self-defense' (1991) 20 *Philosophy and Public Affairs* 283; J Thomson, 'Physician-assisted suicide: two moral arguments' (1999) 109 *Ethics* 497.

explained by the culpable aspects of the wrong. They are not committed to thinking that it is in virtue of the intention that the act is wrong.

Even if (i) does not fail generally, it seems unpromising on most accounts of the moral wrong of ID. This claim will likely appear controversial as some prominent accounts of the wrong do appear to build culpability considerations into its constituent elements. Consider, in this regard, what we can call the 'reasons-based' account of the ID wrong, which endorse the first of the two interpretations of the 'because of' requirement that Altman identified: as indicating 'the reasons that guide the acts of agents'.⁴² One version of this view states that wrongful ID consists (partly) in an agent's causing relative disadvantage by *acting on* an illegitimate reason – a reason which bears some relation to a prohibited ground, such as race. Suppose that a university decides to accept applicants only from local regions because they sincerely believe this reduces the complexity of their admissions procedures. Suppose further that this has a disparate impact on racial minorities. Holmes, followed here by Fishkin, claims that, if this is a wrongful act of discrimination, it can be explained as such by conceiving of the policy (P) as a 'proxy' criterion for race: adopting P-like policies is likely to have a disparate impact on racial minorities.⁴³ Acting on reasons which bear a proxy relation to race-based reasons partakes of the illegitimacy of acting on the race-based reasons. Whatever one thinks of the plausibility of this manoeuvre, it very significantly expands the concept of 'acting on a prohibited/illegitimate reason'. In particular, it seems that one can 'act on an illegitimate reason' on this view even where the ultimate reason why the proxy reason is illegitimate plays no causal role in the explanation of why one took the action.⁴⁴ So even if the university would have adopted the same policy had the racial composition of the surrounding area been entirely different, it still seems to be the case, on this view, that one acts on a prohibited reason here. Further, this account requires no awareness of the relationship the proxy reason bears to the illegitimate reason – in short it requires no subjective or constructive knowledge as to why the proxy reason is indeed a proxy. As such, this divorces the reasons-account from any notion of culpability since it requires no awareness (or objective ability to be aware) of the wrong-making features of the act (namely, its connection to the effects on racial minorities in this example).

Moreau's analysis of the wrong of discrimination might also seem to imply a culpability requirement.⁴⁵ On her view, discrimination is wrongful in so far as it denies a person's equal entitlement to certain deliberative freedoms, in particular freedoms to decide about aspects of one's life without pressure from one's normatively irrelevant characteristics. She describes the wrong as 'tort-like' and, more specifically, as 'akin to the tort of negligence'.⁴⁶ She denies that our equal entitlement is to suffer no adverse effects because of a normatively irrelevant characteristic. Instead, it is rather to suffer no adverse effects because of such characteristics *to the extent this can be reasonably demanded against others who*

⁴² Stanford encyclopedia

⁴³ Holmes, Fishkin, Gardner?

⁴⁴ Cf Essop.

⁴⁵ For an earlier analysis along similar lines, but without the focus on deliberative freedoms in particular: D Reaume, 'Harm and Fault in Discrimination Law: The Transition from Intentional to Adverse Effect Discrimination' (2001) 2 *Theoretical Inquiries in Law* 349

⁴⁶ Moreau, 'Discrimination as Negligence', 130.

cause such effects: “Adverse effect discrimination is not an act or practice that has the effect of excluding someone on the basis of a prohibited ground, but rather an act that excludes someone in circumstances where there is some reasonable way of accommodating this person short of undue hardship”.⁴⁷ Here, then, the primary ID duty is one of reasonable accommodation. When an employer argues that its policy is justified on grounds of business necessity, it does not justify or excuse a wrong, but rather disputes its existence entirely.⁴⁸ While Moreau establishes wrongfulness, nothing here imports a requirement of culpability. An act is unreasonable (and, therefore, wrongful) on this view if it does not instantiate the right balance between the (legitimate) interests of the discriminator and the deliberative freedoms of the victim. But whether the economic interests of the business owner are more important than the deliberative freedoms of the victim is an objective question which holds independently of anyone’s attitudes or knowledge. It is not surprising that Moreau is conscious of the need to allay concerns that her account involves strict liability. It does.⁴⁹

Even expressive accounts which locate the wrongness of (some) acts of ID in the, say, demeaning message communicated by acts of discrimination can accept that what is communicated by an action in a certain community could be wholly unknown to the communicator, who is reasonably unaware of the relevant linguistic or social conventions.⁵⁰

In sum, then, since it seems unlikely that we should believe (i) in relation to ID, (i) gives us no reason to conclude that the law departs from a plausible moral conception of the wrong of ID.

(ii) Culpability as necessary for responsibility for a wrong

What about (ii) (i.e. culpability as necessary for *responsibility* for a wrong)? As a general matter, moral responsibility, at least in one important sense, can diverge from culpability. We can be morally responsible for our wrongful acts even though it would be inappropriate to *blame* us (and even though the act may itself not be *blameworthy*).⁵¹ Consider cases of agents who commit justified wrongs. If, in order to save the lives of passengers, the captain of a ship anchors the ship to a dock owned by another person without that person’s consent during a storm, damaging the dock, the captain wrongs the dock owner, and is responsible for the wrong, but she is in no other way culpable.⁵² (ii) therefore seems untrue.

⁴⁷ Ibid, 134.

⁴⁸ Sometimes, though, Moreau seems to speak of the interference with (certain) deliberative freedoms as itself an infringement of a right without the need any balancing exercise, which suggests the balancing is best understood as a defence to the wrong, rather than a constituent element: see *ibid* 139, cf 140.

⁴⁹ Even if, as Moreau suggests at one point, the causation requirement in ID will involve considerations of foreseeability, this only goes to show that some degree of avoidability might be built into the ID-wrong. But this need not imply that the agent was culpable. See the next section for the divergence between avoidability and culpability.

⁵⁰ Lippert-Rasmussen?

⁵¹ An act might be blameworthy but yet it is inappropriate to blame the agent. See Scanlon, *Meaning Permissibility Blame*

⁵² Cf *Vincent v Lake Erie Transportation Co* 109 Minn 456 (1910).

(iii) Absence of culpability as an excusatory condition

Most, perhaps all, excuses involve denials that the (wrongful) act in question was blameworthy.⁵³ More specifically, excuses concede *responsibility* for the wrong, but deny that the wrong reflects badly, in some sense, upon the agent.⁵⁴ So when a person points to the fact that they committed a wrong because another threatened them with violence, they do not, or need not, deny that the wrong was their doing and that they can be called to account for it, but they do deny that they should be blamed for it. Similarly, if the next time I type 't' on the screen, an earthquake is caused, it would be inappropriate to *blame* me for this – I had no reason whatsoever to believe that this consequence would result. However, that one is excused only blocks an inference from *wrongdoing* to *blameworthiness*. It does not follow that one is relieved of any moral obligations arising out of one's wrong. In other words, only certain kinds of response to one's wrongdoing are ruled out by the absence of culpability. For instance, it seems plausible that one could still be under secondary moral obligations of repair. If this is so, as the law, at any rate, holds, in the case of the justified wrongdoer in tort law (the example of the sea captain), *a fortiori* in the case of the excused wrongdoer. Consequently, since the legal responses to the breaches of the ID duty are not necessarily responses which express blame or which trace their rational source to blaming, the law need not be at variance with morality here. This would only be so if the legal responses to ID necessarily communicated blame in circumstances where they have a genuine excuse.

(iv) Justifying liability in the absence of culpability; Culpability as an aggravating feature of the wrong

Even if there can be moral responsibility without culpability, and even if the absence of culpability does not entail the absence of normative consequences for one's wrongdoing, strict legal liability for ID still calls for justification.

This is for two reasons. First, there is a concern of fairness that agents should not be subject to legal liabilities for wrongs which they cannot avoid committing at the time of acting. If liability is imposed, as it is for ID, without a requirement that it be reasonably foreseeable at the time of A's acting (at the time of breach of the duty) that the 'provision, criterion, or practice' puts (or would put) persons sharing B's characteristic at a disadvantage, then they are subject to liabilities which they may have no reason to believe will arise at the same of acting. Second, the ideal of the rule of law seems also to demand that people be able to assure themselves at the time of acting that their acts will not lead to legal liability.⁵⁵ In short, the ability to avoid the commission of a wrong seems to be required both by considerations of fairness and the ideal of the rule of law.

These are important concerns to which we offer two immediate responses. First, it might be said that a person *can* avoid breaching an ID duty by avoiding entering into one of the legally recognised roles/statuses (such as that of an employer, a landlord and so on) to which the ID duty attaches. So long as these roles are clearly identified and relatively limited in number, there is a degree of avoidability. It is true that liability will attach even in cases where there is no foreseeability of the likelihood that

⁵³ Though some believe that blameworthiness is defined in terms of absence of excuse rather than the other way around. But see A Simester, 'A Disintegrated Theory of Culpability', 183.

⁵⁴ See J Gardner, *Offences and Defences*, ch 6.

⁵⁵ See J Gardner, above, Some rule of law anxieties

the practice in question will produce a disparate impact. But there is general knowledge that one of the abstract risks of setting criteria for the allocation of valuable resources is that the criteria one selects could produce a disparate impact. Although this avoidability comes at immense personal costs (entailed in avoiding otherwise legitimate occupations), and cannot seriously be recommended to anyone, this is still at considerable remove from holding a person responsible for the destruction of a plane which, by some miraculous causal path, happens when they pick up a lettuce in the supermarket. The risk of disparate impact is a known risk which attaches to the setting of criteria for the allocation of resources.

Second, the degree to which avoidability considerations should be taken into account either within the legal definition of the ID duty or as a matter of a defence to the breach of that duty is partly a function of the kind of consequences which attach to the breach. Obviously, if persons who breached the ID duty were subject to *punishment*, there would be a need to provide a much greater degree of avoidability than the current law allows, and thus a much greater role for considerations of culpability within the definition of the wrong or defences to the wrong. But, as we have already observed, the current law does not even provide compensation as of right for unjustified breaches. An employer who innocently breaches their ID duty is likely only to be subject to a recommendation to alter their policy so as to conform to their initial primary duty.⁵⁶

Apart from the possibility of avoidability and the fact that the degree of avoidability required is a function of the severity of the consequences imposed for the breach of a duty, there are three further reasons for believing that the imposition of strict liability for the breach of the ID duty can be justified to persons subject to the duty, all things considered.

First, alongside the ID duty, the law imposes other duties which help to mollify the strict liability nature of the ID duty. S.149(1) of the Equality Act 2010 imposes an obligation on public authorities to 'have due regard to the need to', *inter alia*, 'eliminate discrimination, harassment, victimisation and any other conduct that is prohibited by or under this Act'.⁵⁷ If an authority is to have 'due regard' to this need, this may require, or at least encourage, the authority to gather information about the potential effects of common policies or practices upon protected groups. This helps to establish a body of information about the likely effects of certain practices which others, including private individuals subject to the ID duty, could potentially draw upon in making their own decisions. By improving the state of knowledge about the effects of certain policies on protected groups, this reduces the possibility that persons subject to the ID duty will be caught off guard by its strict liability effect. At the very least, it seems inconsistent with obligation to 'have due regard' not to investigate further incipient signs of disparate impact.⁵⁸ Even if diligent information gathering prior to the adoption of a practice, criterion or policy fails to highlight potential

⁵⁶ See Equality Act section

⁵⁷ S.149(1)(a)

⁵⁸ Cf *The Queen on the Application of Bapio Action Ltd v Royal College of General Practitioners, General Medical Council* [2014] EWHC 1416: 'Section 149 does not permit a person exercising public functions to identify the need to eliminate discrimination in one of the public functions it exercises and then do nothing about it', [31] (Mitting J). See further S Fredman, 'Addressing disparate impact: indirect discrimination and the public sector equality duty' (2014) 43 *Industrial Law Journal* 349.

disparate impact, reliable data should be available *after* such policy has been implemented. Furthermore, assuming that any claimant is likely to seek internal remedies in the first instance before going to the courts, *legal* liability for future breaches is almost always avoidable.

Second, the strictness of the ID duty provides incentives for duty-bearers to invest in information gathering to avoid liability. If the duty were limited to cases of reasonably foreseeable effects, duty-bearers would have less incentives to invest in information since they could avoid liability by showing that the information was not available to the reasonable person. Of course, given the minimal remedial regime, the extra incentive provided is likely to be a relatively weak one.

Third, defining the duty in a culpability-free manner, while taking considerations of culpability into account at the remedial stage, allows the law to provide clear guidance to duty-bearers, while also accommodating considerations of fairness to the latter.⁵⁹ Incorporating, for instance, a reasonableness standard into the definition of the primary duty would cloud the clarity of the law's message. It might be objected here that strict liability nature of the duty could cloud the law's message in a different way. It might lead to the impression that ID is wrongful only in a technical, purely legal, sense.⁶⁰ This concern seems unjustified for two reasons. First, by associating culpable breaches of the ID duty with potential (albeit unlikely) compensatory liability, the law still sends the message that breach of the ID duty is a serious matter. Second, the strict liability nature of the wrong might actually signify how serious the interest in non-discrimination is to the law. In tort law, our most fundamental interests in bodily integrity and liberty are often protected by strict liability wrongs (e.g. in trespass to the person and false imprisonment). The fact that non-culpable infringements of these interests is *wrongful* is a reflection of the law's commitment to stringent protection of these interests.

In sum, then, the absence of culpability in the definition of the primary duty is defensible. Culpability is not required in order for the law to map onto the moral wrong of ID(i). Moral responsibility for a wrong does not entail culpability (ii). To the extent that considerations of avoidability are relevant, there are opportunities to avoid breach of the strict ID duty and, in any event, the level of avoidability required for responsibility varies with the harshness of the consequences attached to breach. Absence of culpability is not formally recognised as an excusatory condition (iv), but this is one way of understanding the general absence of significant normative consequences for the innocent breach. Different levels of culpability may aggravate the wrong and justify different normative consequences (iii).⁶¹

⁵⁹ Cf M Dan Cohen, 'Decision Rules and Conduct Rules: On Acoustic Separation in the Criminal Law' (1983) 97 *Harvard Law Review* 625

⁶⁰ Cf P Shin, 'Liability for Unconscious Discrimination: A Thought-experiment in the Theory of Employment Discrimination Law', 42 (SSRN version) on unconscious bias, suggesting that recognition of liability for unconscious bias discrimination might have a negative effect on the law's message or perceptions thereof.

⁶¹ We will develop this point and the precise role of culpability in the remedial obligations and liabilities of discriminators in future versions. We also hope to consider whether absence of culpability is necessary, even if not sufficient, to excuse or justify *prima facie/pro tanto* indirect discrimination. And if we are really on form, we might think about whether *pro tanto* or *prima facie* is right here.