

Why ‘Indirect Discrimination’ Is a Useful Legal but Not a Useful Moral Concept

KASPER LIPPERT-RASMUSSEN

Aarhus University

Abstract: A policy (practice, act, etc.) indirectly discriminates against a group, G, if, and only if: (1) it does not reflect an objectionable mental state regarding the members of G; (2) it disadvantages members of G; (3) the disadvantages are disproportionate; and (4) G is a socially salient group. I argue that indirect discrimination is not non-instrumentally morally wrong. Clearly, *if* it were, that would be because it harms members of G disproportionately, i.e., in virtue of features (2) and (3). Harming members of a group disproportionately does appear non-instrumentally wrong. But it is not easy to provide a plausible explanation for the kind of harm and disproportionality involved here that vindicates this initial appearance. This does not mean the concept of indirect discrimination should be jettisoned. It was originally a legal concept, and in closing I briefly suggest that in law it plays a valuable role, even if it is not a genuine moral category. Legal prohibition is an unreliable guide to what is morally wrong, but it is not supposed to be that anyway.

Keywords: discrimination, disproportionality, distributive justice, harm, indirect discrimination

JEL Classification: D63, J71

I. INTRODUCTION

Indirect discrimination is not non-instrumentally wrong, i.e., it is not the case that it is wrong independently of its causal consequences—or so I

EJPE’s NOTE: This piece is an invited contribution on the theme of discrimination and inequality.

AUTHOR’S NOTE: I am grateful to the participants in an EJPE roundtable on the Philosophy and Economics of Measuring Discrimination and Inequality, 13 May, 2022; to the participants in a Post Hoc Interventions workshop at Lund University, 6 April, 2022; to the participants in a workshop at the Uehiro Centre for Practical Ethics, 15 June, 2022; and to Lennart Ackermans, Annalisa Costella, Tom Douglas, Mattias Gunnemyr, Binesh Hass, Martin Jönsson, Aksel Sterri, and two anonymous reviewers of this journal for helpful comments.

shall argue in this article.¹ For all I say in this article, direct discrimination might be non-instrumentally wrong and, thus, different from indirect discrimination in this respect, but I will not address the question here. Direct and indirect discrimination are generally distinguished, and generally the distinction is treated as exhaustive and exclusive: it is assumed that, on the basis of a specific feature on which the property of being discriminatory supervenes, all cases of discrimination against a particular group are either direct or indirect, but not both (Lippert-Rasmussen 2013, 40).² Most theorists take the fact that the discriminator's treatment of the discriminatee is driven in some way by, or reflects, an objectionable mental state of his or hers to be a defining feature of direct discrimination. What that mental state is—e.g., an intention to exclude (Altman 2020), feelings of animosity (Arneson 2006), false beliefs about the discriminatee's inferior moral status (Alexander 1992), or an indifference to the interests, or failure to attend to the autonomous choices, of the discriminatees (Eidelson 2015)—and how it is thought to explain the moral wrongness of direct discrimination varies. But for present purposes the crucial thing is that indirect discrimination is defined by the absence of any such mental state.

This absence is not the only way in which indirect discrimination differs from direct discrimination.³ However, it is the most obvious. Given this, it is natural to ask: When a policy (or practice, act, etc.) is not associated with an objectionable mental state, what makes it discriminatory? As the US term for indirect discrimination, 'disparate impact', indicates, the policy's *disadvantaging* of those who are discriminated appears to be

¹ The disproportionate disadvantage for the discriminatee is not a causal consequence of indirect discrimination (see the definition of indirect discrimination below), but a defining feature of it. To say that indirect discrimination is not non-instrumentally wrong means that instances of indirect discrimination might be wrong either in virtue of their defining features (e.g., the disproportionate advantages for the discriminatees that the policy involve in the specific cases), wrong because of their consequences, or not wrong at all.

² This formulation allows that an instance of discrimination could at one and the same time be an instance of direct discrimination against *one* group and an instance of indirect discrimination against *another* group. Setting aside issues pertaining to social salience (see below) a rule requiring employees to be taller than 175 cm could be directly discriminatory against short people and indirectly discriminatory against women. The formulation also allows that a point-based hiring system, which includes the rule 'Deduct 10 points from applicants who are either female or not taller than 175 cm', is both directly discriminatory against women in virtue of the rule's first disjunct and indirectly discriminatory against women in virtue of its second disjunct.

³ Thus direct (unlike indirect) discrimination need not disadvantage discriminatees and is not tied to disproportionality.

crucial here. Hence, for the purposes of this article I define indirect discrimination as follows:⁴

A policy (practice, act, etc.) is *indirectly* discriminatory against a certain group, G, if, and only if:

- (1) it does not reflect an objectionable mental state regarding members of G (for short, the *mental state condition*);
- (2) it disadvantages members of G (the *disadvantage condition*);
- (3) the relevant disadvantages are disproportionate (the *disproportionality condition*);⁵
- (4) and G is a socially salient group (the *social salience condition*; Halldenius 2005, 459).⁶

⁴ The following different, but very similar, characterizations of indirect discrimination have been offered: (a) it occurs when “an act impacts members of a protected group disproportionately, in comparison with its cognate groups” (Khaitan 2018, 32). (b) It pertains to “a certain group, P, if (1) it reflects no bias on the part of the discriminator against members of P on account of them being members of P or any intent to discriminate against members of P (*the no-intention condition*); (2) it unjustifiably disadvantages one group relative to another relevant contrasting group (*the relative disadvantage condition*); and (3) P can be considered as a group in the relevant sense (*the group condition*). All three are necessary to identify cases of indirect discrimination” (Cosette-Lefebvre 2020, 342). (c) “Where a general policy or measure has disproportionately prejudicial effects on a particular group, it is not excluded that this may be considered as discriminatory notwithstanding that it is not specifically aimed or directed at that group” (Shanaghan v. United Kingdom, App. No. 37715/97 Eur. Ct. H.R., para. 129 [2001]; quoted in Altman 2020). (d) “A practice indirectly discriminates against a person, P, on the basis of *t*, if P has *t*, P is disadvantaged by the practice, and although the practice does not explicitly single P out because of *t* or some related trait, *u*, it nevertheless disproportionately disadvantages those who have *t* relative to those who have not” (Moreau 2020, 17). (e) Indirect discrimination occurs where “a facially neutral measure [...] impacts more harshly on one group than on another” (Doyle 2007, 538; see also Fredman 2018; Halldenius 2005, 459; Khaitan 2015, 73–76; Thomsen 2015). As Doyle (2007, 537–538) points out, on the so-called purpose model of indirect discrimination a facially neutral criterion is used to exclude members of a particular group. Such cases are best thought of as cases of (concealed) direct discrimination, so I ignore them here.

⁵ In a legal context, business necessity or a suitable relation to job performance are sometimes taken to rule out indirect discrimination (Griggs v. Duke Power Co., 401 U.S. 424, 431 [1971]). Arguably, business necessity and proper job performance-relation are covered by the disproportionality condition. For a related point about the necessity condition in the context of just war theory, see Hurka (2005, 38).

⁶ Some might adopt a narrower, asymmetric view on which only dominated socially salient groups can be subjected to indirect discrimination (for an overview of anti-subordination in legal theory, see Balkin and Siegel 2003; see also Sullivan 2004). I set aside this view here. If we adopt it, the question arises as to whether a policy that is directed at a non-dominated group and is in all morally relevant respects like an indirectly discriminatory policy against a dominated group (e.g., it disproportionately harms members of the non-dominated group) is any worse, morally speaking, than the otherwise comparable indirectly discriminatory policy. ‘Otherwise comparable’ could be understood in such a way that the same disadvantage in one sense of ‘disadvantage’ imposed on a member of a dominated group and on a member of a dominating group (e.g., not being hired) typically involves very different degrees of disadvantage in the relevant sense. For

The mental state condition certainly seems to capture a feature that we are latching on to when we see a difference between direct and indirect discrimination. As I have indicated, it is also a generalization of accounts of indirect discrimination in the literature.⁷ The disadvantage condition seems incontrovertible.⁸ In a case of direct discrimination it might make sense to complain about being subjected to discrimination *qua* being a member of G even if, for some reason, one is not harmed by being discriminated against. But the notion that one could benefit from being indirectly discriminated *against* (or indeed be unaffected by indirect discrimination) makes no sense. If anything, such cases are better described as situations in which either there is no discrimination at all or, if the members of G benefit, there is indirect discrimination in *favor* of G.

The disproportionality condition, likewise, is plausible. As noted by Hellman (2018, 108), not just any policy that disadvantages members of G will indirectly discriminate against G (let alone be morally wrong for this or other reasons).⁹ For one thing, it may be the case that all the alternatives to the policy in question impose greater disadvantages with greater moral significance on other groups, in which case complaining about the practice on grounds of indirect discrimination would appear to be misplaced.¹⁰ The disproportionality condition is designed to immunize accounts of indirect discrimination, including the one given in (1) to (4),

example, losing one of very few options for being hired versus losing one of very many options for being hired. On this understanding, we may be able to explain why, say, white, middle-aged, heterosexual men are much less likely to experience indirect discrimination than members of minority groups, even if there is no conceptual barrier to their having disproportionate disadvantages imposed on them.

⁷ I set aside here the complication of policies that do not reflect objectionable mental states but pick out all (or only) members of the group of discriminatees for disadvantageous treatment and do not pick out any members of the group of non-discriminatees for disadvantageous treatment. Plausibly such policies are classified as directly discriminatory. However, this complication makes no difference to my arguments below, so I set it aside.

⁸ I set aside how indirect discrimination should be defined in relation to merely possible people—an issue that might be quite relevant in relation to abortion policies, etc.

⁹ I do not address what the relevant measure of (dis)advantage is in this article. My arguments are neutral regarding the proper metric of (dis)advantage.

¹⁰ It might be suggested that such a complaint makes sense on the grounds that the relevant non-instrumental wrongness thesis merely says that indirect discrimination is *pro tanto* wrong, not wrong all things considered (see below). One rejoinder to this suggestion notes that, typically, when people assess whether the disproportionality condition is satisfied, they assume that only relatively direct effects (whatever 'direct' means here) count for the purposes of assessing proportionality. On that assumption, and given that remoter effects count in relation to whether something is morally wrong all things considered, we should reject the relevant suggestion.

against this problem. Clearly, that leaves us with the task of specifying when harms are disproportionate. I return to this in section IV.

Finally, the social salience condition. A lot of policies that are not considered discriminatory nevertheless impose disproportionate disadvantages on groups of people. The groups, however, are—in a sense needing to be explained—artificial: an example is the group of people who are disproportionately disadvantaged by the policy in question. The social salience condition addresses this issue. The key idea is that a group of people can be subjected to indirect discrimination only if they form a socially salient group, where a group is socially salient if, and only if, “perceived membership of it is important to the structure of social interactions across a wide range of social contexts” (Lippert-Rasmussen 2013, 30). Men and women, African Americans and European Americans, disabled and non-disabled people are socially salient groups in the US. Americans living in a postal district whose number contains two identical numerals are not. The social salience condition finds strong support in that almost all complaints about indirect discrimination concern socially salient groups, and that, typically, in cases where socially salient groups are subjected to indirect harms in a way that does not reflect objectionable mental states, indirect discrimination is detected.¹¹

In this article, I want to use the definition of indirect discrimination in (1) to (4) as a springboard from which to explore the concept of indirect discrimination and ask why this form of discrimination is considered morally wrong. By ‘morally wrong’ I mean (mostly without stating this explicitly) *pro tanto* morally wrong, i.e., that there is something about indirect discrimination which makes it to some degree morally wrong. An act or practice can be *pro tanto* morally wrong, yet not morally wrong all things considered. An instance of indirect discrimination might also have various good-making features such that all things considered it is not wrong to engage in it. Acts and practices can also be *prima facie* wrong, yet not *pro tanto* wrong. Thus, the fact that an act amounts to indirect discrimination might be a defeasible reason to think that it is wrong for some reason or other even if we know that this reason cannot be that it is a case of indirect discrimination. Hence, the denial that indirect discrimination is *pro tanto* wrong, leaves it open to me to say that it is *prima facie* wrong in the indicated way. Section II argues that if we accept the social salience condition, we must also accept that indirect discrimination

¹¹ See, however Cosette-Lefebvre (2020, 351–357) for a critique of social salience-focused definitions of discrimination.

is not a *distinct* wrong—in other words, cases of indirect discrimination are morally no worse than cases where comparably disproportionate harm is done to a socially 'insalient' group of people.¹² (From here onwards I use this handy expression to refer to groups without lacking salience.) That leaves it open that indirect discrimination is a wrong, but not distinctively so. Section III explores the disadvantage condition in that light. It considers different accounts of the conditions under which indirect discrimination harms the discriminatees with a view to determining whether harming them in one of these ways might be non-instrumentally wrong. Section IV turns to the disproportionality condition, which is underspecified in the above definition. I argue that friends of the view that indirect discrimination is non-instrumentally wrong face a dilemma. Either we understand disproportionality in a non-revisionist way, or we do not. If the former, then indirect discrimination cannot be non-instrumentally wrong. If the latter, it can—at least, this seems a genuine possibility—but then, of course, we are saddled with a revisionist view of indirect discrimination.

I am not the first theorist to defend the claim that indirect discrimination is not non-instrumentally wrong.¹³ The main novelties of this article lie elsewhere. First, it offers several new arguments for this claim that are based on a relatively detailed analysis—a second novelty—of group disadvantage and the disproportionality conditions. In section V it also offers an account of why the concept of indirect discrimination is useful despite the arguments of the previous section. The view that indirect discrimination is not non-instrumentally wrong might be taken to mean that we need not be concerned about that form of discrimination. However, this inference should be questioned. In section V, I explain why, in general, there may well be good moral reasons for being concerned, from the point of view of law, about policies (practices, acts, etc.) that are not

¹² There is another sense in which the wrong of indirect discrimination can be said to be distinct: in virtue of being different from the wrong of direct discrimination. Although this is not the distinctiveness I reject in section II, Sophia Moreau (2020) does deny that the wrong of indirect discrimination is distinctive in this sense. On her view, both direct and indirect discrimination are wrong because they involve a failure to treat the discriminatees as moral equals.

¹³ I defended the same view in Lippert-Rasmussen (2015) and Eidelson (2015, 39–68) defends a similar view. Also, theorists—say, libertarians—who think that discrimination is wrong only in a quite narrow range of cases take the same view as I do on whether indirect discrimination is non-instrumentally wrong (e.g., see Cavanagh 2002). Recently, Moreau (2020) has defended the claim that direct and indirect discrimination are both non-instrumentally wrong, and for the very same reason, i.e., that they involve failing to treat the discriminatees as equals.

wrong in themselves. It argues that these reasons may well be applicable to indirect discrimination. Section VI concludes.

II. SOCIAL SALIENCE

Consider a policy which is disproportionately harmful to a particular group of people who do not form a socially salient group.¹⁴ Compare that policy to one that is exactly alike except that the people harmed do form a socially salient group. My contention is that these two cases cannot differ, morally speaking. Why not?

The most obvious answer is that a given harm is not made smaller simply by being imposed in a particular way—in this case by being perpetrated in a way that visits the harm on members of a particular (socially insalient) group. Hence, insofar as it is differences in the magnitude of harm that make a difference to wrongness, then the social salience or otherwise of the group cannot make the two cases morally different. It is true that harm is unlikely to be the only wrong-making feature. Arguably, objectionable mental states such as animosity (see section I) are another. However, because our concern here is with indirect discrimination and relevantly similar policies, the two cases I am comparing do not differ in those terms.

It might be replied that something must have gone wrong here. Imposing a certain amount of disproportionate harm on, say, a vulnerable, socially salient minority seems to be morally worse than imposing the same amount of harm on a group of people who do not form a socially salient group and who are not similarly vulnerable. I concede that the two cases do indeed look morally different. However, this is entirely consistent with the claim that I am making here. Note, first, that a group of people can be vulnerable even if they do not form a socially salient group. Hence, if we assume that the two cases differ in terms of vulnerability, we are comparing cases of disproportionate harm to two groups of people that differ only in the fact that one group is socially salient, and the other is not.

Second, certain vulnerabilities are closely tied to membership of a socially salient group. For example, the members of a socially salient group, G, may be vulnerable to negative stereotypes about members of G in ways that it is difficult to see members of a socially insalient group being. Presumably, groups that are not socially salient typically do not attract

¹⁴ Recall my definition of a socially salient group in section I.

stereotyping (Beeghly 2021). But while this is a morally relevant difference, it is not a difference which by definition—as opposed to: as a general matter—overlaps with the distinction between groups that are socially salient and groups that are not. Accordingly, it cannot show that social salience is non-instrumentally morally significant, and thus it cannot show that imposing disproportionate disadvantages on a socially salient group is morally worse than imposing those disadvantages on a group which is not.¹⁵ I conclude that while indirect discrimination might be wrong, if it is, it is not distinctively wrong—by which I mean ‘wrong in a way that otherwise comparable instances not amounting to indirect discrimination are, by contrast, not wrong’. The question then is: *Is* indirect discrimination non-instrumentally wrong?¹⁶ My response to this question will occupy the next two sections.

III. GROUP DISADVANTAGE

The disadvantage condition says that indirectly discriminatory policies disadvantage members of the groups being subjected to discrimination. This condition is rarely specified in the respects I want to consider, but it can be understood in several ways. Consider the classic court case involving indirect discrimination, *Griggs v. Duke Power Co.* (401 U.S. 424 [1971]).¹⁷ Here, is the disadvantage condition satisfied if the company’s

¹⁵ Something similar is true of expressive harms. While it might be true that, virtually always, only members of socially salient groups are subjected to expressive harms *qua* their group membership, in principle members of insalient groups could suffer expressive harms too. For further arguments against it being non-instrumentally morally relevant whether a group of people disadvantaged by a certain policy form a socially salient group, see Thomsen (2013).

¹⁶ On the desert-accommodating, prioritarian account of the wrongness of discrimination that I proposed in Lippert-Rasmussen (2013), it is both the case that the wrong of indirect discrimination is not distinct from the wrong of direct discrimination and that the wrong of discrimination, whether direct or indirect, is not distinct from otherwise comparable, but non-discriminatory, policies etc. Non-distinct moral wrongs can, in my sense, be serious moral wrongs. Note also that while section II defeats some suggestions as to what might make indirect discrimination distinctively wrong, it does not demonstrate that there cannot be any feature of indirect discrimination which makes it distinctively, non-instrumentally wrong.

¹⁷ This 1971 US Supreme Court verdict implied that a hiring policy can be illegal when it is “fair in form, but discriminatory in operation” (*Griggs v. Duke Power Co.*, 401 U.S. 424, 431 [1971]), i.e., indirectly discriminatory. The employer, Duke Power, had “instituted requirements of high school education and satisfactory scores in an aptitude test as a condition of employment or transfer. The same test was applied to all candidates, but because black applicants had long received education in segregated schools, both requirements operated to disqualify black applicants at a substantially higher rate than whites” (Fredman 2011, 178). Since the requirements went beyond what was needed to ensure satisfactory levels of performance, the company was ordered by the court to abolish them.

seniority policy disadvantages *its own African American employees*? Or is it satisfied if that policy (or the set of similar policies operated by several companies including Duke Power) disadvantaged African Americans *in general*? It probably resulted in both, even if its effects on the Duke Power employees were much more significant than its largely marginal effects on African Americans generally. From our point of view the crucial question is about which of following four views articulates the correct understanding of the disadvantage condition (for convenience, I consider discriminatory *policy* from this point onwards):

A policy satisfies the disadvantage condition if, and only if:

Local View: it disadvantages a relevant subset of G, e.g., employees in the relevant company.

Global View: it disadvantages members of G in general, e.g., members of G on average or each individual member of G.

Broad View: it satisfies *either* the local *or* the global view.

Restrictive View: it satisfies *both* the local *and* the global view.

I think that insofar as indirect discrimination is non-instrumentally wrong *both* local *and* global disadvantages must matter. In my view, an act or practice can qualify as wrongful, indirect discrimination either because of the disadvantages it imposes on a subset of the relevant group members or because of the disadvantages it imposes on group members in general even if it does not involve both forms of disadvantage. If so, the Restrictive View is too restrictive.¹⁸ Consider the following case, designed to test the supposition that global effects are irrelevant to whether something amounts to discrimination:

Hypothetical Griggs v. Duke Power Co.: This case is like the actual one except that while Duke Power's promotion rules disadvantage its African American employees, for some weird reason this benefits African Americans in general.¹⁹

¹⁸ It is too restrictive because, presumably, a policy may be indirectly discriminatory if it disproportionately disadvantages all the relevant company's employees even if it does not disadvantage members of the relevant minority in general.

¹⁹ It is standard in many philosophical disciplines to appeal to our intuitions about hypothetical examples like the present one (in fact, often examples much more outlandish than this one are presented). However, some theorists are sceptical of whether we can learn anything from the intuitions we have about hypothetical cases (see also Moreau

Plausibly, for some sufficiently small size of the relevant local harms and some sufficiently large size of the global benefits for other members of the relevant group, the connected policy would not be *pro tanto* wrong even if it were indirectly discriminatory. Granting this, and granting also that indirect discrimination is non-instrumentally wrong, it follows that the disadvantage condition does not apply only to harms to the relevant subset of the group of people being subjected to indirect discrimination. Perhaps in many cases indirect discrimination only disadvantages members of the relevant group locally, and we may wish to say that this suffices to satisfy the disadvantage condition. However, that does not mean that global disadvantages are irrelevant. It also follows that on a non-revisionist understanding of it, indirect discrimination is not non-instrumentally wrong, since the local disadvantages could be substantially outweighed, morally, by the global advantages to the group to which the discriminatees belong.

Whether we accept the Broad View, as I have just suggested we should, or instead accept one of the three alternative interpretations of the disadvantage condition, there is a further question about what the disadvantage condition says, namely: What exactly is it to disadvantage the relevant (subset of the) group in question?

There are simple cases where this question is not pressing, because *all* members of the relevant group are disadvantaged by the indirectly discriminatory policy. However, real-life cases of indirect discrimination—especially if the disadvantages condition applies globally—are rarely simple. Typically, some members of a group subjected to indirect discrimination are disadvantaged and others are not affected—and some might even benefit all things considered. Hence, we need to ask when a group is disadvantaged all things considered. The most obvious answer is:

The Average View: Members of a group are disadvantaged by a certain policy if, and only if, the policy disadvantages members of the group on average.

While the Average View is a natural view to take, and although it has intuitive results in a range of cases, it cannot be right if indirect discrimination is non-instrumentally wrong—at least, not if we adopt a widely

2020, 29). I do not have the space to discuss this methodological issue. For a view which is moderately sympathetic to the use of thought experiments in applied ethics, see Walsh (2011).

accepted account of the aggregation of harms and benefits. Consider a case where the great majority of group members benefit marginally from a policy while a minority suffer very significant harm from it—so significant, in fact, that people who oppose unrestricted aggregation would say that it is morally impermissible to impose that very significant loss on a few people to benefit a much greater number of people even though the aggregate benefits are much greater than the aggregate harms (cf. Scanlon 1998, 229–241; see also Voorhoeve 2014).²⁰ Here we might be inclined to say that the policy mistakenly alleged to be indirectly discriminatory is wrong because it disproportionately harms members of a certain group *even if* members benefit on average and *even if* most members of the group benefit.²¹ However, we might think instead that if indirect discrimination names a particular moral wrong, it is better to classify it as a case of indirect discrimination. With this attitude, we might adopt:

The Anti-Unrestricted Aggregation View: Members of a group are disadvantaged by a policy if: (i) some of its members are severely disadvantaged by the policy; (ii) the aggregate disadvantages imposed on these members is smaller than the aggregate benefits conferred on other members of the group; (iii) but the individual benefits conferred on these members are normatively insignificant relative to the harms to the former.

Suppose we are anti-aggregationists. That is, we believe that there are some trivial benefits which, by comparison with some significant harms, are such that, however large the number of people receiving the former and however small the number of people suffering the latter, it is morally impermissible to impose significant harms of that type on the smaller number of people even if that brings about a (much) greater sum of

²⁰ This claim does not rest on any assumptions about how well off the recipients of the marginal benefits and the recipients of the significant harms are—either absolutely or relatively—before being benefited or harmed in the relevant ways. The claim is consistent with the view that such assumptions make a difference to whether a benefit or harm of a certain magnitude (i.e., a certain amount of prudential value) counts as marginal or significant from a moral point of view.

²¹ This case also enables us to see that another intuitively appealing view—the *Numbers View*, which says that members of a group are disproportionately disadvantaged only if a (super)majority of the group's members are disadvantaged—is false. Thomsen (2015, 321) describes a variant of the Numbers View: a policy disadvantages P-persons if “a greater proportion of P-persons than of P-persons are negatively affected by the treatment, or [...] P-persons are more severely affected than P-persons by the treatment”. Assuming indirect discrimination is wrong non-instrumentally, as Thomsen (2015, 322) in effect acknowledges, this view is problematic for a reason like that defeating the Numbers View.

benefits. If that is the case, and if indirect discrimination is non-instrumentally wrong, we must reject the Average View and accept the Anti-Unrestricted Aggregation View. At any rate, not doing so would render the fact that a certain practice is not indirectly discriminatory less significant, since the absence of indirect discrimination would then be compatible with the impermissible imposition of severe harms on a small subset of group members. Even worse, there could be cases of indirect discrimination which are not non-instrumentally wrong, because the only alternative involves the imposition of severe, but aggregately smaller, harm on a tiny subset of group members.

Even setting aside the issue of unrestricted aggregation, a further reason to reject the Average View is that it is insensitive to the distribution of harms and benefits across members of a group. If that distribution is morally significant, it will be a factor in when we want to say that the disadvantage condition is satisfied. In the interest of simplicity, let us focus on a prioritarianism—the view that benefiting “people matters more the worse off these people are” (Parfit 1991,19). On this view, two policies imposing the same amount of overall harm on the same number of members of a particular group could differ in that one is indirectly discriminatory, and the other is not. In one case most of the harm is done to worse off members of the group, whereas in the other most of the harm is done to better off members. Assuming the policy also delivers certain benefits to members of the groups such that, on the Average View, none of them counts as disadvantaging members of the group, we might think that such a case is a reason to reject the Average View, since, given its terrible distributive profile, the policy in question here does, in the normatively relevant sense, disadvantage members of the group. Instead, we might adopt:

The Distribution-Sensitive View: Members of a group are disadvantaged by a certain policy if, and only if: the policy disadvantages members of the group on average; and the relevant disadvantages are weighted according to how badly off the recipient of the disadvantage is—the worse off she is, the greater weight the disadvantage has.

Again, it seems we cannot consistently accept that distributive justice involves distribution-sensitivity, the Average View, and the view that indirect discrimination is non-instrumentally wrong. If we accept the Average View, we might have to classify a policy as indirectly discriminatory and yet accept that it is not even *pro tanto* wrong because any alternative

would involve a morally worse distribution of harms, e.g., one where harms of a similar size fall mostly on people who are already worse off.²² Equally, we would have to recognize the possibility of policies involving a bad distributive profile which we could not classify as indirectly discriminatory because on average the group members benefit. Yet we would consider these policies unjust. Again, this would render it less significant that a certain policy is not indirectly discriminatory. After all, that sought-after status would then be entirely consistent with the severe harms imposed on worse off members of the group by the policy rendering it impermissible. Hence, the consideration regarding non-aggregation and distribution suggests that on a non-revisionistic understanding of the group disadvantage condition indirect discrimination is not non-instrumentally wrong.

Finally, the views discussed so far imply that we can assess whether something amounts to indirect discrimination without knowing anything about the causal history of the disadvantages in question. Arguably, this is not the way standard cases of indirect discrimination have generally been understood. Take *Griggs v. Duke Power Co.* again. There, the company's promotion rules harmed African Americans as a result of historical (direct) discrimination against them, since that meant that, as the rules were applied, no African Americans had the seniority needed for promotion. The promotion rules compounded previous unjust harms. Had the case not involved this feature—had it been untrue, in other words, that the distribution of seniority when the rules were applied was itself expressive of discrimination—then, arguably, the disadvantage condition would not have been satisfied.²³ The view underpinning this asymmetric assessment might be:

The Anti-Compounding View: Members of a group are disadvantaged by a certain policy in the relevant sense if, and only if, the policy

²² I assume that the mere fact that there is an individual who is worse off under a certain distribution than she would be under some alternative distribution is not enough to make it *pro tanto* wrong (as opposed to being in one respect worse). An option I am drawn to here is that it is only if the individual is worse off under a particular distribution relative to the morally required distribution that her being worse off is *pro tanto* wrong.

²³ This is consistent with the claim that almost any disadvantage imposed by a policy on standard victims of discrimination compounds an injustice. My case here simply rests on the modest contention that it is *possible* for the disadvantage not to be a compounding of a prior unjust disadvantage. For a real-life case which arguably did not involve compounded injustice, yet was seen by courts as a case of indirect discrimination, see Hellman (2018, 121).

disadvantages members of the group, and those disadvantages are causally dependent on prior unjust discrimination against this group.²⁴

On this view, a policy does not indirectly discriminate against a group if it disadvantages if the disadvantages are causally independent of prior unjust discrimination against the group.²⁵ However, indirect discrimination might take place if the disadvantages in question are so dependent. If we think that indirect discrimination is wrong when it compounds prior injustice (Hellman 2018, 106), and if indirect discrimination is non-instrumentally wrong, it seems to follow that the Anti-Compounding View of disadvantage is crucial to a correct understanding of the disadvantage condition. At least, if disadvantages are understood in terms of anti-compounding—something which Hellman (2018, 121) acknowledges is somewhat revisionist relative to standard legal notions of indirect discrimination—it seems to follow that indirect discrimination is non-instrumentally wrong. If, however, indirect discrimination is wrong when it compounds prior injustice, and if we reject the Anti-Compounding View of disadvantage, it follows that indirect discrimination is not non-instrumentally wrong.²⁶ This follows because, given these two assumptions, there will be cases of indirect discrimination that do not compound injustice, since the disadvantage condition can be satisfied in the absence of any compounding of injustice.

To sum up, the disadvantage condition is much more complex than it looks. Specifically, if indirect discrimination is non-instrumentally morally wrong, we must, in a revisionist way, take that condition to pertain to both local and global disadvantages. We then face a question about how to weigh, say, local disadvantages against global benefits and vice versa. Also, it seems we must reject the Average View. Finally, we might even have to take a Restrictive View of what counts as a disadvantage in the

²⁴ The Anti-Compounding View raises the same issues of interpretation that the Average, the Anti-Unrestricted Aggregation, and the Distribution-Sensitive Views all address. The causal dependency condition is related to a similar, but distinct, causal dependency condition proposed in Lippert-Rasmussen's (2013, 71) definition of indirect discrimination. The latter refers to past or present direct discrimination, while the present condition refers to prior unjust discrimination, thus including unjust indirect discrimination.

²⁵ A slightly broader view includes injustices other than unjust discrimination.

²⁶ The present point does not require us to accept Hellman's (2018) view that indirect discrimination is wrong because it compounds injustice, and here I am simply using that view as a pointer to a certain interpretation of the group disadvantage condition. In passing, however, I should mention that I am not persuaded by Hellman's account for reasons explained in Eidelson (2021) and Lippert-Rasmussen (forthcoming); for a reply to Eidelson, see Hellman (2021).

relevant sense—that is, insist that it is only disadvantages compounding prior injustice that so count. I have not shown that there is no possible interpretation of 'disadvantage' on which indirect discrimination can be non-instrumentally wrong, but I have shown, I hope, that what such a view would look like is significantly more complicated than is normally assumed. I also hope to have shown that adopting a view of disadvantage on which it is possible to claim that indirect discrimination is non-instrumentally wrong requires one to take a distinctly revisionist view of several dimensions of the cases that are to be counted as genuine cases of indirect discrimination. The dilemma this generates has not, as far as I am aware, received attention in the literature so far.

IV. DISPROPORTIONALITY

Suppose we have come up with a convincing answer to the issues discussed in the previous section. That still leaves us with the challenge of explaining the disproportionality condition. This challenge has two aspects. First, if a disadvantage is disproportionate, it is disproportionate relative to something else. But what, exactly, is the thing we are to compare the disadvantage to the indirect discriminatees with in order to determine whether its imposition is disproportionate? Call this the issue of the proper *comparanda*. Second, once we have fixed what the proper comparanda are, we need to determine what it is for them to be related in a way exhibiting disproportionality. For example, is it sufficient for the disadvantage to indirect discriminatees to be greater to some degree than the corresponding benefits to others or must it be significantly greater?²⁷ Call this the issue of *quantity*. In this section, I address both issues with a view to determining whether indirect discrimination can be a non-instrumentally wrong.

Let me start with the issue of the proper comparanda, drawing on (Lippert-Rasmussen 2013). The two most interesting views here are:

²⁷ In the previous section, I explored issues about unrestricted aggregation, and distribution, and whether the relevant disadvantage compounds a prior injustice. Similar issues come up in relation to the relevant positive comparanda, as it were. For example, do the benefits to others fall on people who are very badly off, or on people who are very well off? Suppose that a policy harms well-off women significantly but benefits badly off men slightly. On a distribution-sensitive view of proportionality, this policy might not satisfy a distribution-sensitive version of the proportionality condition. For simplicity, I set these matters aside in this section and assume that we should simply compare the size of the disadvantages to indirect discriminatees with the size of the benefits to others.

The Group Comparison View: A group, G_1 , is disproportionately disadvantaged by a policy if, and only if, the inequality between G_1 and G_2 (a group with which it is to be compared) is greater with the policy than it would be in some relevant alternative situation without the policy.²⁸

The Advantages Comparison View: A group, G_1 , is disproportionately disadvantaged by a policy if, and only if, the gap between the advantages it would enjoy under that policy relative to the greater advantages that G_1 would enjoy under some relevant alternative policy is disproportionate relative to the gap between the advantages some group, G_2 , with which it is to be compared would enjoy under that policy relative to the smaller advantages G_2 would enjoy in some relevant alternative situation without it (Lippert-Rasmussen 2015).²⁹

To see how the two views differ, consider a situation where we must choose between two policies, *Equal* and *More*. Both men and women will be best off under *More*, which however involves inequality in men's favour. Under *Equal* policy there will be strict equality between men and women, but both groups will be worse off than they would have been under *More*. Here, the Group Comparison View implies that women are disproportionately disadvantaged by *More*—the inequality between men and women is greater under *More* than it is under *Equal*. The Advantages Comparison View need not have this implication, because neither men nor women are disadvantaged under *More* in comparison with how well-off they would be under *Equal*.

On the whole, those alleging indirect discrimination either do not distinguish between Group Comparison and Advantages Comparison. If they do, they are likely to assume that, in practice, a policy that is disproportionate in one sense will be disproportionate in the other as well.

²⁸ Khaitan (2018), Cosette-Lefebvre (2020), and Doyle's (2007) definitions of indirect discrimination (see note 4) all seem to imply the Group Comparison View.

²⁹ A more formal definition of the two views is the following: Let $A(g, p)$ be the advantage that group g enjoys under the policy p . A group, G_1 , is disproportionately disadvantaged (in comparison with another group, G_2) by a policy α (as compared with an alternative policy β) if, and only if: *The Group Comparison View:* $A(G_2, \alpha) - A(G_1, \alpha) > |A(G_1, \beta) - A(G_2, \beta)|$. *The Advantages Comparison View:* $A(G_1, \beta) - A(G_1, \alpha) >_* |A(G_2, \alpha) - A(G_2, \beta)|$ where '>*' means not just 'greater than' but 'disproportionately greater than'. I owe this definition to an anonymous reviewer. Note also that because the definiendum is a policy that disproportionately disadvantages a certain group, it does not render the definition of the Advantages Comparison View problematically circular that the definiens refers to a disproportionate gap between differences in outcome for one group relative to the differences in outcome for another group. In short, different notions of disproportionality appear in the definiendum and the definiens. I thank the editors for pointing out the need to clarify this matter.

However, as *More* and *Equal* show, the two views can be pried apart. Accordingly, to say whether indirect discrimination is non-instrumentally wrong, we need to consider both explications of the comparanda of disproportionality. There are troubles ahead whichever view we side with.

Take, first, the Group Comparison View. On this view, unless you subscribe to distributive egalitarianism in a group-focused version, you cannot regard indirect discrimination as non-instrumentally wrong. But many of those who write about justice reject distributive egalitarianism in the light of the so-called levelling down objection (or for other reasons, e.g., because they think that only inequalities between individuals have moral significance). That objection asks us to consider a situation of equality in which no one is better off in any respect than they would have been in an alternative situation without equality, and then invites us to agree that the equal situation is in no way better than the latter unequal situation (Parfit 1991). People who are impressed with the levelling down objection yet started out with egalitarian sympathies typically adopt prioritarianism. But in the prioritarian perspective, a policy that is indirectly discriminatory, if the disproportionality condition is interpreted in terms of the Group Comparison View, might be one that brings about the best outcome and thus is not non-instrumentally wrong.³⁰

The Advantages Comparison View avoids this problem, but it creates another: the problem of revisionism. For it implies that some cases that most of us would consider to be clear cases of indirect discrimination do not satisfy the disproportionality condition. Thus consider Khaitan's (2018, 31) description of indirect discrimination as "an apparently neutral practice or policy which puts members of a protected group (say, women) at a disproportionate disadvantage compared with members of a cognate group (say, men)". The idea here, and elsewhere when people see indirect discrimination, is clearly that, roughly, we can simply compare the proportion of, say, male to female employees to determine whether indirect discrimination has taken place.³¹ However, on the Advantages Comparison View it is possible for a policy to result in, say, a greater proportion of male (or female) employees even if it is not indirectly discriminatory.

³⁰ Even on a strict egalitarian view, disproportionality as presently understood need not be a non-instrumentally wrong-making feature. This is so because the currency of the disadvantages involved in the relevantly indirectly discriminatory rule (e.g., proportion of women employees) typically is not the currency of egalitarian justice (e.g., access to advantage [Cohen 1989]).

³¹ 'Roughly' because the differential representation must somehow result from the policies (practices, acts, etc.) of the indirectly discriminating agent and be disproportionate relative to the benefit obtained through the policy (Khaitan 2018, 41).

On this view, indirect discrimination might be non-instrumentally wrong even if we reject distributive egalitarianism, but only if we think differently about the conditions under which a policy is indirectly discriminatory and re-classify various seemingly clear-cut cases of indirect discrimination as cases that are free of that form of discrimination.

Turning to the second issue, about quantity, there are two ways to understand the disproportionality condition here: in a moralized or non-moralized way. On the first, by definition, whether a certain disadvantage is disproportionate entails something about its moral qualities, whereas on a non-moralized understanding this is not the case. On what is probably the simplest moralized analysis, a disadvantage is morally disproportionate relative to the corresponding benefits, if, in the light of the relationship between the two, it is morally unjustified to impose it. On the moralized interpretation, if a disadvantage is disproportionate, it is morally wrong by definition to impose it on the victims of indirect discrimination. On a very straightforward non-moralized analysis, a disadvantage to the victims of indirect discrimination is disproportionate if the harm that falls on them is greater than the advantages the discrimination gives to others. Here, though this might be the case, it does not follow from the fact that a disadvantage is disproportionate that it is morally unjustified to impose it. Where the disadvantage *is* unjustified, that is so not because indirect discrimination has occurred, but because some independent moral principle has been infringed.

Neither understanding of proportionality is unproblematic for our purposes. If we embrace the moralized conception, we cannot say that a policy is made morally wrong by being indirectly discriminatory. Rather, one reason that policy *is* indirect discrimination is that it is morally wrong.³² However, many of those who write about discrimination think that the fact that something is indirect discrimination *makes* it wrong. Hence, they need to opt for the non-moralized view.

On the non-moralized view, however, discrimination simply falls out of the picture in the following sense: it is the fact that indirect discrimination involves the imposition of disadvantages on one group of people—where these are in a suitable descriptive sense disproportionate relative

³² By way of analogy, consider murder. Murder is a moralized concept in that, roughly, murder is a morally wrongful type of killing. But even if something's being murder entails that it is morally wrongful, what *makes* it morally wrongful cannot be that it is murder. Rather, part of what makes the action in question murder is that it is morally wrongful. For a helpful discussion of the distinction between what makes something discrimination and what makes it wrong, see Ishida (2021).

to the benefits to others—that makes it wrong. However, as I argued in section II, there is no reason to think the wrongness of such an imposition is tied narrowly to indirect discrimination, since, among other things, it could also arise in connection with harms to members of a socially insalient group. Moreover, it is unclear that this kind of disadvantage-imposition is wrong in general. Indirect discriminators typically adopt all sorts of policies—e.g., about where to build facilities, what sort of research to invest in, which goods to produce, and so on—that benefit some people and harm others (Hellman 2018, 108). Like decisions about recruitment and promotion, all these policies could be assessed in terms of disproportionality. However, few of us believe that they should be so assessed, and certainly not according to the standards of proportionality used in typical indirect discrimination cases.³³ But if indirect discrimination is morally wrong because of its disproportionality, so must be decisions about where to build facilities, what sort of research to invest in, and which goods to produce.³⁴

Can we say that this shows many more decisions are wrong because they are indirectly discriminatory than we suppose? That might be a sensible response for those who think indirect discrimination names a particular moral wrong. After all, it is hard to see why, say, hiring policies could be *pro tanto* morally wrong in virtue of their disparate effects on minorities, whereas investment policies with similar disparate effects are not. Ultimately, however, many will be reluctant to go down this particular route given the radical expansion of policies which, potentially at least, could then be shown to amount to wrongful indirect discrimination. They will instead prefer to respond to the present line of argument by rejecting the view that indirect discrimination is non-instrumentally wrong and then explaining why—for pragmatic reasons connected, for example, with the successful operation of the law—we nevertheless have reason to treat indirectly discriminatory hiring decisions differently from many other decisions resulting in disproportionate disadvantage being imposed on

³³ Judging by their behaviour, most people think it is permissible in daily life to make decisions about how to spend one's money which are strongly disproportionate because they involve tiny benefits for them (dining out at a fancy restaurant) and not preventing much greater harms to others (by providing them with medicine through a donation to Oxfam).

³⁴ Indeed, it is an exception that decisions are considered in the light of whether they involve disproportionate harms to different socially salient groups. Probably, in relation to many decisions, many of us would say that it is the relevant agent's prerogative to select an option that is disproportionately harmful according to the standards of proportionality used in typical indirect discrimination cases.

certain groups. Something along the latter lines is what I shall briefly pursue in the next section.

V. WHAT IS THE FUNCTION OF THE CONCEPT 'INDIRECT DISCRIMINATION'?

Let me stress again that the alternative to thinking that indirect discrimination is non-instrumentally wrong is *not* to think that many of the policies that are rightly considered indirectly discriminatory should be permitted by law. For we might think that indirect discrimination is not non-instrumentally wrong, but that nevertheless it is a good thing that law forbids indirect discrimination. In fact, there are two good reasons for thinking this is indeed so. The first is that indirect discrimination law is an instrument we can use to secure certain socially desirable ends, and no clearly better alternative is available. The second is that if we treat indirect discrimination in this way, we can set aside many of the puzzles above as irrelevant, or as relevant but no threat to the desirability of indirect discrimination law.

Let me start with how indirect discrimination law is a useful instrument even if the discriminatory policies it outlaws are not non-instrumentally morally wrong. Many Americans and foreign observers believed that the ban of direct discrimination introduced in the era of the civil rights movements would result in greater racial equality and integration. This did not happen. One reason is the difficulty proving that an employer has discriminatory intent. A related reason is the limited incentive to avoid discrimination that a discrimination law only covering direct discrimination creates. Indirect discrimination law avoids this problem (Collins and Khaitan 2018, 25-26). Once such a law is in place, just about any significant underrepresentation of a protected group will oblige the employer to prove that the skewed staffing is not due to discrimination. Shifting the burden of proof from complainants to employers in this way has no doubt considerably strengthened the employer's incentive to recruit in ways that promote equality and racial integration.³⁵ No instrument is

³⁵ An anonymous reviewer suggested that the notion of indirect discrimination can play a similar useful role in applied ethics in that significant underrepresentation of a group in a way that amounts to indirect discrimination against the group is a *prima facie* reason to think that the policies producing such underrepresentation are morally wrong. Hence, the concept of indirect discrimination is a useful heuristic, as it were, in applied ethics because showing that certain policies are indirectly discriminatory shifts the burden of proof from those who think the policies are unjust to those who think they are not. I accept this point. When I submit that indirect discrimination is not a useful moral concept I have in mind the purposes of normative theory—to develop an account of

perfect, of course, and no doubt indirect discrimination law has negative side-effects. However, the crucial point here is that there is *an* important reason in its favor—its contribution to racial equality and integration—even if indirect discrimination is not non-instrumentally morally wrong.

So much for the outcome-based case for indirect discrimination law. How does viewing indirect discrimination law as a useful instrument rather than a regulator of policies, practices and acts that are inherently wrong help with the problems described in the three previous sections? I shall take the three conditions of indirect discrimination that I have examined above in turn.

Consider, first, the social salience condition. I argued that there is no moral difference non-instrumentally between a policy that disproportionately disadvantages a socially salient group and an otherwise similar policy that disproportionately disadvantages a socially insalient group. If we regard indirect discrimination law as a useful instrument mitigating specific injustices, this observation is neither here nor there. For, first, the fact that two wrongs are morally similar does not show that an instrument that serves well in preventing one of them will effectively prevent the other. Second, even if there is no non-instrumental moral difference between a policy that disadvantages a socially salient group and an otherwise similar policy that disadvantages a group lacking that salience, it may still be the case that there are many more policies of the former kind than the latter. Perhaps, for example, it takes less for a disadvantage to be disproportionate when it affects underprivileged groups than it does when it falls on privileged people—and if so, it may well be sensible to prohibit indirect discrimination against certain 'protected groups' without prohibiting any imposition of disproportionate disadvantage on any group (see note 6).

Next, I argued that the disadvantage condition needs to be specified in several key dimensions. In particular, we need to know whether the disadvantage is local or global, and whether it depends on (a relevant subset of) group averages or something more complicated. In the context of law, these issues appear much less pressing, because, as indicated in section III, in real life, and for the protected groups typically covered by indirect discrimination law, it is likely that the relevant groups will be disadvantaged whichever of the specifications we embrace. No doubt, there will be exceptions, and in those cases discrimination law will not serve

wrong/right making features—not the purposes of applied ethics, which are *inter alia* to assess the moral qualities of specific policies etc.

our moral aims optimally. However, a suboptimal instrument can be the best instrument available.

Consider, finally, the disproportionality condition. In section IV, I discussed the comparanda and quantity of disproportionality. If we treat indirect discrimination law as a useful legal tool, we might respond to the first issue in the way I just responded to the issue of specifying the relevant sense of disadvantage. Thus, we might say that, by and large, when a protected group is disadvantaged by a policy on the Group Comparison View, it will be disadvantaged by that policy on the Advantages Comparison View, and vice versa, and that accordingly we can regard indirect discrimination law as a useful instrument whichever view we embrace. Responding to the issue of quantity, we can simply observe that indirect discrimination law typically errs on the side of caution: arguably, whichever plausible view of proportionality we adopt, indirect discrimination law typically outlaws policies only if they are significantly disproportionate and is tolerant of policies if they are less significantly so even when—perhaps, ideally—law ought not to be tolerant of policies that involve moderately disproportionate disadvantages being imposed on certain groups. Accordingly, if our aim is to defend existing indirect discrimination law (not to embrace an expansion of the scope of existing indirect discrimination law) failure to provide a satisfactory answer to the question of quantity is not a big problem.³⁶

The notion of indirect discrimination, then, may well be useful in a legal context, because anti-indirect discrimination law could promote or otherwise serve important moral values. I have sketched an argument to this effect, but not actually given it. In addition, I have briefly shown why, if we think about indirect discrimination in this way, the objections to the view that indirect discrimination, understood in a non-revisionist way, is non-instrumentally wrong considered in other parts of the paper (or suitably adjusted variants of them) need not worry us insofar as we adopt the perspective on indirect discrimination I have explored and supported.

VI. CONCLUSION

I have argued that, on reflection, we can see that there are real challenges in understanding indirect discrimination as something that it is non-instrumentally wrong. More briefly, and pointing in the direction I think

³⁶ This is not to say that the issue of quantity never gives rise to problems for the proposed law-centric view of indirect discrimination. For one thing, I have not shown that anti-indirect discrimination law is never overinclusive.

that this leads, I have tried to show that it makes good sense to regard indirect discrimination as a useful legal concept, nevertheless. Legal moralists think that if something is morally wrong, we have reason to make it unlawful. What we might call 'moral legalists' are inclined to think that if something is unlawful it is *pro tanto* morally wrong (and would remain so even if it were not prohibited by law). I believe both outlooks are false, and in section V I have offered some reasons for thinking that the legal status of indirect discrimination is a poor guide to the moral status of indirect discrimination.

While I am not the first theorist defending this view about the moral status of indirect discrimination (e.g., see Eidelson 2015, 68), many writers have taken a different view (e.g., Moreau 2020, 187; Moreau 2018).³⁷ Hence, it might be useful to end this article by highlighting three issues that are *not* part of the disagreement between those—myself included—who think that indirect discrimination is not non-instrumentally morally wrong and those who do. First, everything I have said above is consistent with saying that *direct* discrimination is non-instrumentally morally wrong.³⁸ Second, the same goes for the view that most (apparent) instances of indirect discrimination are morally objectionable. For example, it may be that most such instances are really cases of direct discrimination (Eidelson 2015, 6), or it may be that either on their own or together with other similar policies the relevant instances result, or have resulted, in systematic, unjust disadvantages for certain groups. Finally, and partly in light of the previous point, nothing in this article motivates a strongly revisionist stance regarding indirect discrimination law.³⁹

REFERENCES

- Alexander, Larry. 1992. "What Makes Wrongful Discrimination Wrong? Biases, Preferences, Stereotypes, and Proxies." *University of Pennsylvania Law Review* 141 (1): 149-219.
- Altman, Andrew. 2020. "Discrimination." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Article published February 1, 2011; last modified April 20, 2020. <https://plato.stanford.edu/entries/discrimination>.
- Arneson, Richard. 2006. "What Is Wrongful Discrimination?" *San Diego Law Review* 43 (4): 775-807.

³⁷ For an overview of a range of views about the role of indirect discrimination law, see (Collins and Khaitan 2018, 25-29).

³⁸ I do not accept that it is (Lippert-Rasmussen 2013). However, my reasons for this are independent of the case I have made here in connection with indirect discrimination.

- Balkin, Jack M., and Reva B. Siegel. 2003. "The American Civil Rights Tradition: Anticlassification or Antisubordination?" *University of Miami Law Review* 58: 9–33.
- Beeghly, Erin. 2021. "What Is Wrong with Stereotypes? The Falsity Hypothesis." *Social Theory and Practice* 47 (1): 33–61.
- Cavanagh, Matt. 2002. *Against Equality of Opportunity*. Oxford: Clarendon Press.
- Cohen, Gerald A. 1989. "The Currency of Egalitarian Justice." *Ethics* 99 (4): 906–944.
- Collins, Hugh, and Tarunabh Khaitan. 2018. "Indirect Discrimination Law: Controversies and Critical Questions." In *Foundations of Indirect Discrimination Law*, edited by Hugh Collins and Tarunabh Khaitan, 1–30. Oxford: Hart Publishing.
- Cosette-Lefebvre, Hugo. 2020. "Direct and Indirect Discrimination: A Defense of the Disparate Impact Model", *Public Affairs Quarterly* 34 (4): 340–367.
- Doyle, Oran. 2007. "Discrimination, Indirect Discrimination, and Autonomy." *Oxford Journal of Legal Studies* 27 (3): 537–553.
- Eidelson, Ben. 2015. *Discrimination and Disrespect*. New York, NY: Oxford University Press.
- Eidelson, Benjamin. 2021. "Patterned Inequality, Compounding Injustice, and Algorithmic Prediction." *American Journal of Law and Equality* 1 (1): 252–276.
- Fredman, Sandra. 2011. *Discrimination Law*. 2nd edition. New York, NY: Oxford University Press.
- Fredman, Sandra. 2018. "Direct and Indirect Discrimination: Is There Still a Divide?" In *Foundations of Indirect Discrimination Law*, edited by Hugh Collins and Tarunabh Khaitan, 31–56. Oxford: Hart Publishing.
- Halldenius, Lena. 2005. "Dissecting 'Discrimination'." *Cambridge Quarterly of Healthcare Ethics* 14 (4): 455–463.
- Hellman, Deborah. 2018. "Indirect Discrimination and the Duty to Avoid Compounding Injustice." In *Foundations of Indirect Discrimination Law*, edited by Hugh Collins and Tarunabh Khaitan, 105–122. Oxford: Hart Publishing.
- Hellman, Deborah. 2021. "Personal Responsibility in an Unjust World: A Reply to Eidelson." *American Journal of Law and Equality* 1 (1): 277–285.
- Hurka, Tom. 2005. "Proportionality in the Morality of War." *Philosophy & Public Affairs* 33 (1): 35–66.
- Ishida, Shu. 2021. "What Makes Discrimination Morally Wrong?" *Theoria* 87: 483–499.
- Khaitan, Tarunabh. 2015. *A Theory of Discrimination Law*. New York, NY: Oxford University Press.
- Khaitan, Tarunabh. 2018. "Indirect Discrimination." In *The Routledge Handbook of the Ethics of Discrimination*, edited by Kasper Lippert-Rasmussen, 30–41. New York, NY: Routledge.
- Lippert-Rasmussen, Kasper. 2013. *Born Free and Equal? A Philosophical Inquiry into the Nature of Discrimination*. New York, NY: Oxford University Press.
- Lippert-Rasmussen, Kasper. 2015. "Indirect Discrimination Is Not Necessarily Unjust", *Journal of Practical Ethics* 2 (2): 33–57.
- Lippert-Rasmussen, Kasper. Forthcoming. "Is There a Duty not to Compound Injustice?" *Law and Philosophy*.
- Moreau, Sophia. 2018. "The Moral Seriousness of Indirect Discrimination." In *Foundations of Indirect Discrimination Law*, edited by Hugh Collins and Tarunabh Khaitan, 123–148. Oxford: Hart Publishing.
- Moreau, Sophia. 2020. *Faces of Inequality*. New York, NY: Oxford University Press.

- Scanlon, Thomas. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Sullivan, Charles A. 2004. "The World Turned Upside Down? Disparate Impact Claims by White Males." *Northwestern University Law Review* 98 (4): 1505–1565.
- Thomsen, Frej K. 2013. "But Some Groups Are More Equal Than Others." *Social Theory and Practice* 39: 120–146.
- Thomsen, Frej K. 2015. "Stealing Bread and Sleeping beneath Bridges—Indirect Discrimination as Disadvantageous Equal Treatment." *Moral Philosophy and Politics* 2 (2): 299–327.
- Voorhoeve, Alex. 2014. "How Should We Aggregate Competing Claims?" *Ethics* 125 (1): 64–87.
- Walsh, Adrian. 2011. "A Moderate Defense of the Use of Thought Experiments in Applied Ethics." *Ethical Theory and Moral Practice* 14: 467–481.

Kasper Lippert-Rasmussen is Professor of philosophy at the department of Political Science at Aarhus University. He works in political philosophy, and has conducted extensive research on discrimination, egalitarianism, and fairness. Presently, he is the director of the Centre for Experimental-Philosophical Studies of Discrimination funded by the Danish National Research Foundation.

Contact e-mail: <lippert@ps.au.dk>