

Sophia Moreau
Draft Paper – September 2006

What is Discrimination?

1.

Relatively little philosophical work has been done on the nature of discrimination.¹ But discrimination raises a number of philosophically interesting issues. When we identify an action or a rule as discriminatory, we are claiming not just that it draws a distinction between people or results in some being treated differently from others, but also that this differentiation is in certain respects unfair. What exactly are these respects: in what does the unfairness consist? Why does it seem particularly objectionable, from a moral standpoint, to subject others to this type of unfair treatment? And is it ever justifiable to impose legal sanctions on discriminatory actions, even though this would appear to interfere substantially with citizens' freedom to choose whom they contract with and what sorts of reasons they act upon when they do so?

One reason why these questions have not been pursued in detail in philosophical discussions may be that it is natural to treat intentional discrimination as the paradigmatic case of discrimination, and there is a common view of intentional discrimination that provides answers

¹ For exceptions, see Matt Cavanagh, *Against Equality of Opportunity* (Oxford: OUP, 2002) and Larry Alexander, "What Makes Wrongful Discrimination Wrong?" (2002) 141 U. Pa. L. Rev. 149. More philosophical work has been done on affirmative action, which I take to raise somewhat different issues. See, for instance, Thomas Nagel, "Equal Treatment and Compensatory Discrimination," (1973) 2 *Philosophy and Public Affairs* 348; Thomas E. Hill, Jr., "The Message of Affirmative Action" (1991) 8.1 *Social Philosophy and Policy* 108; and Elizabeth Anderson, "Integration, Affirmative Action, and Strict Scrutiny," *NYU Law Review*, 77 (2002): 1195-1271.

to the above questions that seem relatively straightforward, and so philosophically uninteresting.² Intentional discrimination –for instance, refusing someone entry to your bar because he is black, or denying someone a job because she is a woman-- occurs when actions or rules are adopted with the purpose of excluding some people because they possess a certain trait. Of course, this general description also fits many innocuous acts of exclusion, such as denying a job to those who lack the skills required by the job. According to a common view of intentional discrimination, what distinguishes these innocuous exclusions from those that constitute intentional discrimination is the motive for which the exclusion is undertaken: either the agent's sole motive is to cause harm to those excluded, or, although the agent is acting in furtherance of some other purpose, he believes it acceptable to exclude these people in order to further this purpose because he sees them as inherently inferior beings, who simply deserve less. So intentional discrimination, on this common view, involves deliberately excluding someone on the basis of a certain trait, either with the sole aim of harming her (we can call this "malice") or in the belief that it is acceptable to harm her in the service of some other goal because her possession of this trait renders her inferior as a human being (we can call this "prejudice"). It is not particularly difficult to explain why such actions are morally abhorrent. Most moral views that treat the content of a person's motives as relevant to the moral status of her actions imply that it is morally wrong (or in some other way amounts to a moral failing) to act out of malice –that is, to set out to harm others, simply for the sake of causing harm to them. And most also imply that it is similarly objectionable to act from the belief that others are inherently inferior to

² Those whose discussions presuppose what I am calling the "common view" of intentional discrimination include Cavanagh and Alexander, *supra* note 1, and also John Gardner, "Liberals and Unlawful Discrimination", (1989) 9 Oxford J. Legal Stud. 1; and Richard Primus, "Equal Protection and Disparate Impact: Round Three" (2003) 117.2 Harv. L. Rev. 493

oneself, as human beings –that is, to act from what I have called prejudice. Moreover, it is natural to suppose that these morally objectionable motives are also the source of our justification for legally prohibiting discrimination. Most countries that prohibit discrimination by one citizen against another do so only in certain special contexts, contexts where the discriminator has chosen to offer certain goods to the public such as employment, accommodation, or goods and services. It seems plausible to suggest that when we choose to take on such roles in the public domain, then we must, like the state, treat everyone as though they were of equal worth. And it can seem obvious that, whatever else this means, it must mean at least that we cannot act with the sole purpose of causing harm to some of them, or act on the belief that some are inherently inferior to others.

One of the subsidiary aims of this paper is to question this understanding of intentional discrimination. I shall suggest that not all cases of intentional discrimination involve a belief in the inferiority of others or a desire to cause harm to them. More importantly, I shall argue that even in cases that do involve such objectionable mental states, the mental states themselves do not play any role in justifying the legal prohibition on such discrimination; though, to the extent that they cause greater suffering for the victim, they may justify requiring the agent to pay additional compensatory damages to her. Since the focus of my paper will be on the legal prohibition of discrimination, I shall have less to say about why discrimination is morally objectionable; but it is an implication of my argument that there is a more complicated story to tell than the one sketched above.

The main aim of the paper, however, is to offer an analysis of a different form of discrimination. This is what is known as “adverse effect” or “disparate impact” discrimination.³ This form of discrimination can occur in the absence of a belief about the inferiority of those who are excluded or a desire to harm them, and indeed, even in the absence of any exclusionary purpose at all on the part of the discriminator. If an action or rule has the effect of excluding certain people because of their possession of a certain trait and this trait is a recognised “ground” of discrimination, then that action or rule may amount to adverse effect discrimination. The word “may” is necessary because there is a proviso here, which different countries interpret in slightly different ways. I shall work with the Canadian interpretation, as it seems one of the clearest and most promising. On this interpretation, the proviso reads that such actions amount to adverse effect discrimination if, and only if, there is some reasonable means of accommodating the excluded persons which falls short of imposing “undue hardship” on the agent. So, for instance, the Royal Canadian Mounted Police’s former policy of requiring the traditional “Stetson” hat to be worn at all ceremonies came to be regarded in the 1990’s as adverse effect discrimination against practising Sikhs, whose religion requires them to wear a turban. It was, of course, not the purpose of the policy to exclude Sikh officers. But the policy had the effect of excluding them from ceremonies because of their religion. Because the only consequences for the RCMP of accommodating these officers were aesthetic and historical ones (it would mar the RCMP’s uniformity of appearance and constitute a break with tradition), and

³ It is referred to as “adverse effect discrimination” in Canada, “disparate impact discrimination” in the United States, and “indirect discrimination” in the U.K. I shall use the Canadian terminology throughout my paper.

neither of these were seen to involve important freedoms on the part of the RCMP, it was felt that no undue hardship would be imposed by requiring accommodation.

To see the structure of adverse effect discrimination and to understand how the defence of undue hardship works, it may help to consider another case and three variants on it.⁴ Suppose a blind customer with a guide dog tries to have lunch at a local cafe. She sits down and waits for service. But the service never comes: no waiter approaches her. On the first variant of this case, let us suppose this is because they simply do not see her: perhaps there is a large tree in front of her table, whose gorgeous foliage obscures her. In this case, she cannot complain of discrimination because the fact that she was not served did not result from any action or policy on the part of the waiting staff. So although they may be faulted on other grounds –for instance, for having failed to be properly attentive to potential clientele-- they cannot be said to have performed an action that resulted in her exclusion, much less her exclusion on the basis of a prohibited ground. On the second variant of this case, let us suppose that they do see her. But the restaurant has a firm policy that no dog is to be allowed on the premises because dogs lower the tone of the establishment. So they ignore her, knowing that sooner or later she will take her dog elsewhere. This second variant does, likely, constitute adverse effect discrimination: it seems unreasonable of the restaurant to refuse to serve this client when doing so would not materially harm the premises or substantially interfere with the running of the restaurant. It is true that, in this second variant, the waiting staff have not deliberately set out to exclude the

⁴ The case I discuss here is loosely based on the facts of *Douglas Parisian v. Hermes Restaurant Ltd.* (1987), 9 CHRR D/4756 (Manitoba Court of Queen's Bench) and *Commission des droits de la personne et des droits de la jeunesse c. 9107-9194 Quebec inc. (Restaurant Jing Hua)*, [2005] J.T.D. P.Q. no. 24, though neither of these cases involved the religious dimension that I have added to variant 3.

blind client: it is her dog that they want out of the restaurant, and if she were willing to leave him by the entrance, they would happily serve her. But because the guide dog's constant presence is required by her, given her blindness, their action has the effect of excluding her on the basis of a prohibited ground of discrimination. Finally, consider a third variant on this case. Suppose that the restaurant staff have a religious objection to approaching the dog: on their interpretation of their religion, a brush against a dog would render them ritually impure and hence unable to complete the next set of prayers without cleansing themselves. Suppose it is a small restaurant –perhaps with only one waiter— so that for the waiter to have to cleanse himself frequently as a result of encounters with a guide dog would pose substantial difficulties for the operation of the restaurant. In this third variant on the case, we might think that accommodating this client really would constitute an undue hardship for the restaurant. And so we might conclude that this third variant is not, after all, a case of adverse effect discrimination: there is, in these circumstances, no reasonable accommodation possible that falls short of undue hardship because the only ways of accommodating this client are ones that impose an unacceptably large burden on the restaurant.

Adverse effect discrimination, then, is any act or rule that has the effect of excluding someone on the basis of a trait that is a prohibited ground of discrimination, in circumstances where there is some reasonable way of accommodating this person that does not impose undue hardship on the agent. So understood, adverse effect discrimination can seem philosophically puzzling in at least two respects, respects that might lead us to question whether we can ever

justifiably use state coercion against someone for actions or rules that amount to this form of discrimination.

The first respect concerns the apparent absence of any sort of fault or wrongdoing on the part of the agent.⁵ If adverse effect discrimination is just a matter of the unfortunate side-effects that some actions and rules have on particular individuals because of characteristics such as their race, gender, or religion, then it is unclear how we could ever justifiably treat it as due to wrongdoing on the part of the person who performed the action or adopted the rule. What, exactly, did they do wrong? There is, of course, an obvious but unhelpful sense in which any such agent acts wrongly: when he violates legal prohibitions on adverse effect discrimination, he contravenes positive law, and in this sense he commits a wrongdoing. But this is not the sense of “wrongdoing” that I am concerned with here. I am concerned with whether adverse effect discrimination involves wrongdoing in some non-tautological sense –some sense that would *justify* the legal prohibition on it. And it can seem as though it does not. Was it, for instance, really the fault of the RCMP that Sikh officers were excluded from participation in RCMP ceremonies? Was it, likewise, due to some wrongdoing on the part of the waiters that, in variant 2 of my restaurant case, the blind client was not served? All that the waiters did was notice a dog and decide that they did not want this dog in their restaurant. It is true that it was this decision of theirs that caused the blind client to be excluded in a way that others were not excluded, just as it was the RCMP’s policy that caused the officers to face a burden that other

⁵ Note that I am using “fault” here –and throughout the paper-- in the legal, rather than the moral sense. So, as I am using the term, it does not imply that the agent is morally blameworthy.

officers did not face. But most actions and policies impose an unequal burden on someone or other. Why should we treat these cases –cases in which people are excluded because of some trait that we treat as a prohibited ground, even though there are reasonable accommodations available—as ones where the burden amounts not just to an unfortunate side-effect of the discriminator’s action but to wrongdoing on his part? One may want to answer: grounds of discrimination mark out groups that have historically been excluded and underprivileged, and it is important to redistribute opportunities and resources to these groups. One way of doing so is by legally coercing employers and service providers to make reasonable accommodations for the needs of these groups. But although this may help to explain why it is socially beneficial to have laws placing these burdens on the shoulders of employers and service providers, it does not tell us how these laws could be just, in the absence of some fault or wrongdoing on the part of these individuals.

If we cannot explain why adverse effect discrimination involves wrongdoing on the part of the discriminator, then we face at least two problems. Firstly, it becomes much harder to justify the claim that the agent of adverse effect discrimination can be legally coerced into shouldering the cost of whatever measures are required to accommodate those whom his acts have excluded. And most regimes of private sector anti-discrimination law assume that this is justifiable: when a case of discrimination is identified, it is the agent who is required to cover the full costs of accommodating the victim. If the exclusion of the victim is not due to his wrongdoing, however, then why should the costs be transferred onto his shoulders? Why not instead have the costs covered by the public at large through a general taxation scheme, so that

the burden is spread evenly across all of us? Secondly, if adverse effect discrimination does not involve wrongdoing, then there is no sense in which the victim of adverse effect discrimination can legitimately complain that she has personally been wronged by the agent of it. And this runs counter to a further important fact about the way in which we conceptualize private sector discrimination, and in particular the way in which we conceptualise the kind of wrongdoing that it involves. Up until now, I have simply been speaking of “wrongdoing” on the part of the discriminator. But we do not treat adverse effect discrimination as an instance of what Cardozo might have called “wrongdoing in the air”.⁶ Rather, we think that it is the kind of wrongdoing that generates a personal claim on the part of the victim. We believe that victims of discrimination have a specific grievance against the discriminator: they have not just suffered a loss as a result of his misconduct, but have rather been treated by him in a way that amounts to a personal wrong against them. And this belief, too, is central to private sector discrimination law. Most systems of private sector discrimination law use a “complaint model”: that is, they rely on individual complainants to instigate legal proceedings against alleged discriminators and the process is conducted throughout as an investigation into whether the complainant has herself been wronged by the discriminator. Although carriage of the complainant’s case is sometimes given to a legal body that is acting in the wider public interest, such as a human rights commission, it is still the case that this body’s primary role in the proceedings is to bring forward and resolve what is viewed as a personal dispute between the complainant and the alleged discriminator. Moreover, many jurisdictions allow for special “dignitary” damages to be paid to

⁶ In *Palsgraf v. Long Island Railroad* 248 N.Y. 339, Cardozo spoke of negligence as a tort that should be conceptualised not as “negligence in the air”, but rather as negligence toward a particular person or group of persons, in respect of a certain action. I am suggesting here that this is also how we normally conceptualise adverse effect discrimination.

the victims of discrimination, over and above the costs of redressing whatever inconveniences or material losses the victims have suffered. These special damages seem simply to be a symbolic acknowledgment of the fact that these individuals have been wronged.

Hence, one puzzle raised by adverse effect discrimination is how it could be seen as involving wrongdoing, in such a way as both to support our belief that those who discriminate can legitimately be coerced into shouldering the cost of accommodating those whom they have excluded, and to support our intuition that the kind of wrongdoing at issue involves a personal wrong against the victims of discrimination. But there is also a second puzzling feature of this form of discrimination. This is that it seems to have very little in common with intentional discrimination –or at least, with the common view of intentional discrimination that I sketched above. Indeed, they seem to be two quite different phenomena. Deliberately excluding someone because one desires to harm her or believes that this treatment is appropriate given her inferiority is a very different thing from acting in a way that has the undesired side-effect of disadvantaging her. The distinction is not quite directly parallel to Holmes' famous distinction between kicking the dog and tripping over it.⁷ For tripping over the dog is not, under any description, an intentional action; whereas the agents who bring about adverse effect discrimination are at least acting intentionally under some description. They are just not acting intentionally under the description that makes it seem discriminatory. So the distinction between the two forms of discrimination seems closer to the distinction between setting out to kick the dog because you want to harm him or because you think this is deserved or acceptable given his inferior status,

⁷ Oliver Wendell Holmes, *The Common Law* (Transaction Publishers, 2005) 7.

and either wittingly or unwittingly stepping on the dog because he is lying in the middle of the sidewalk and you are in a hurry. But this is still a significant difference. Moreover, if intentional discrimination is legally objectionable for the reasons we examined earlier --namely, the motives for which it is done-- whereas adverse effect discrimination is legally objectionable because of its effects, then the legal prohibitions on the two sorts of discrimination have two completely different sources. This seems odd, given that we speak of these two forms of discrimination as if they were just that: two forms of the same phenomenon, that should be objectionable for at least some of the same reasons.

There are several responses to these puzzles that seem initially attractive, but that I think we have reason to reject.⁸ One is to deny that there really is any objectionable form of discrimination other than intentional discrimination. Adverse effect discrimination, insofar as it merits the name “discrimination”, is just what New Zealand’s early *Human Rights Commission Act* referred to as “discrimination by subterfuge”: that is, it is a malicious or prejudicial motive hidden behind a facially neutral act or rule.⁹ So adverse effect discrimination really does involve wrongdoing and really is the same phenomenon as intentional discrimination.

⁸ My aim here is not to offer an exhaustive list of the alternatives to my own view, but only to discuss a few plausible alternatives. I have chosen these ones partly because they seem plausible; but also because each fails to capture some of the basic features of adverse effect discrimination. So by seeing where they go wrong, we can sharpen our view of the phenomenon that we are trying to account for.

⁹ *Human Rights Commission Act, 1977, s.27 (N.Z.)*. This legislation was subsequently replaced by the *Human Rights Act, 1993 (N.Z.)*, which eliminated the concept of “discrimination by subterfuge” and recognised disparate impact discrimination as a distinctive form of discrimination.

This response seems false to our experience of cases of adverse effect discrimination. Most of them would be mischaracterized, and the genuinely difficult issues that they raise, buried, if they were seen as cases of concealed malice or prejudice. To return to my earlier example of the RCMP's ceremonial uniform, it seems highly implausible to suggest that the committee responsible for setting this uniform policy was really hostile to members of religions that require certain headwear or assumed that such people were inferior and could legitimately be excluded from ceremonial occasions. Most likely, they initially did not think about the impact on these individuals at all; and when it was brought to their attention, they felt it was regrettable but more important to preserve RCMP tradition and uniformity of appearance. So even though they were indeed aware of this impact at a later stage, it was no part of their purpose, and they did not think it was deserved. Many cases of adverse effect discrimination have this pattern: initially, the act or rule is chosen not out of prejudice or a sense of superiority, but simply without any awareness of its impact on certain individuals. And when this impact is brought to the agent's attention, he regrets it but feels that other important reasons outweigh the reasons for changing it. So the effects continue. The question that adverse effect discrimination raises is whether we can make sense of *this* phenomenon as fault-based and continuous with whatever is objectionable about intentional discrimination; and it is no answer to this problem to suggest that it does not exist.

But perhaps a more subtle version of this first response is in order. One might argue, not that there is no such phenomenon as adverse effect discrimination, but simply that because it does not appear to involve wrongdoing, we cannot legitimately prohibit it by law. So whatever

we may or may not say about adverse effect discrimination from a moral standpoint, we cannot say that it is justifiably prohibited by law (or at least, in such a way as to require the agent cover the costs of accommodating the victim). Insofar as we are to make sense of prohibitions on adverse effect discrimination as legitimate laws, we must re-interpret them as prophylactic measures designed to ferret out instances of covert intentional discrimination.

This is a view currently popular among some American legal scholars.¹⁰ They argue that it can be difficult for victims of discrimination to prove a malicious or prejudicial motive, especially in cases where a rule or action is facially neutral. So by requiring proof only of burdensome effects on the victim because of her possession of a trait marked out by a forbidden ground of discrimination, we eliminate the evidentiary difficulties posed by intentional discrimination and we cast a net that is wide enough to catch cases of intentional discrimination that might otherwise be missed. But this net must not be mistaken for one that is designed to catch some other form of discrimination. On this view, then, adverse effect discrimination functions as a kind of *res ipsa loquitur*: just as the latter tort doctrine allows us to treat the defendant's act in the circumstances as "speaking for itself" and obviating the need for any further proof of negligence on the defendant's part, so this way of understanding adverse effect discrimination assumes that certain sorts of objectionable consequences speak for themselves and render proof of ill motives unnecessary.

¹⁰ See, for instance, Christine Jolls, Antidiscrimination and Accommodation, (2001) 115 Harvard Law Review 643; A. Morris, "On the Normative Foundations of Indirect Discrimination Law: Understanding the Competing Models of Discrimination as Aristotelian Forms of Justice" (1995) 15 *OJLS* 199; and George Rutherglen, "Disparate Impact Under Title VII: An Objective Theory of Discrimination" (1987) 73 Va. L. Rev. 1297. See also Justice Stephens' reasoning in *Alexander v. Sandoval*, 121 S. Ct. 1511, 1530 n.13 (2001).

It may be that this response offers us the only possible way of making sense of prohibitions on adverse effect discriminations as legitimate law. If we were to accept this response, however, we would have to abandon many of the current goals of discrimination law. We look to discrimination law to help combat oppression; to uncover and reform institutional structures that disadvantage already stigmatized minority groups; and most broadly, as the *Canadian Human Rights Act* states, to give effect to the principle that “all individuals should have an opportunity equal with other individuals to make for themselves the lives that they are able and wish to have and to have their needs accommodated.”¹¹ This is not, of course, to say that we need (or should) take this expansive statement literally. But it is to say that discrimination law, at least in most countries, has a more radical, reformist agenda than simply ferreting out exclusions motivated by malice or prejudice. So before abandoning this agenda, we should try to see whether it is indeed possible to make sense of adverse effect discrimination as a form of discrimination that does not require malice or prejudice on the part of the agent.

The last response I want to canvass before laying out my own view is one that questions an assumption that has been underlying both my presentation of the problem and the two responses to it that I have already examined. This is the assumption that in order for adverse effect discrimination to be legitimately prohibited by law, we must make sense of it as involving wrongdoing. One might argue that this is mistaken. John Gardner has suggested, for instance, that those individuals and institutions to whom private sector discrimination law applies –for

¹¹ R.S 1985, c. H-6, s. 2.

instance, employers, service-providers, accommodation-providers—are really just “agents of distributive justice”.¹² Their positions, he argues, give them access “to some of the most important opportunities that modern life presents”. Consequently, they are well-placed to redistribute these valuable opportunities from “the opportunity-advantaged” to “the opportunity-disadvantaged” in accordance with whatever principles of distributive justice a society has chosen. Adverse effect discrimination is simply an act or rule that fails to divide up the available opportunities in the manner that accords with society’s chosen principles of distributive justice. So on this view, although adverse effect discrimination is a legal wrong in the positivist sense that it is forbidden by actual laws (and indeed Gardner himself repeatedly refers to it as “wrongdoing”), it does not depend on the perpetrator having in any deeper sense committed wrongdoing, or on his having personally wronged the victim. All that it depends on is his having been a poor or incompetent agent of distributive justice. He has, we might say, failed to aid people in circumstances where they need it, when society’s chosen principles of distributive justice require that the state provide, or find someone to provide, that aid.

One problem with this response is that, at least as it stands, it does not explain why it is fair to people such as employers and service-providers to treat them as “agents of distributive justice”. It may be socially valuable to do so; and it may, as Gardner suggests, help the state to discharge its obligations of distributive justice. But it does not follow from the fact that the *state* has an obligation of distributive justice to bring about a particular state of affairs that this thereby

¹² Gardner, “Discrimination as Injustice” (1996) 16.3 OJLS 353 at 363; subsequent quotations are also from this page.

becomes the obligation of each private citizen who happens to be in a convenient position to help discharge it. More has to be shown, in order to explain why the state can justifiably treat the need for reasonable accommodation as the agent's responsibility, to be paid for by his private funds. Moreover, this account of prohibitions on adverse effect discrimination sits uncomfortably with the rest of Anglo-American private law, which does not normally assume that the state can justifiably order individuals to take positive steps to assist others. Most common law jurisdictions are wary of imposing even a narrower and more urgent positive duty – the duty to rescue in “easy emergency” cases, where the victim's life is at risk and she could be rescued at little or no cost to the rescuer. If it seems problematic to require individuals to take positive steps to aid others in cases where a life is at risk and rescue would impose little or no cost on the rescuer, how much more problematic is it to demand that individuals take positive steps to accommodate others' needs in cases where no one's life is at stake and the costs of accommodating others can be significant?

A further problem with Gardner's way of understanding adverse effect discrimination is that it simply admits that the second of our puzzles –namely, how to make sense of adverse effect discrimination as the same sort of phenomenon as intentional discrimination—is intractable. It implies that there are two radically different forms of discrimination, which are legally prohibited for very different reasons. Gardner himself explicitly acknowledges this, but holds that it is not a problem: on his view, discrimination can be of different types, each of which can be wrong for a number of different reasons. Although this may be right, it seems preferable to try first to work out a view of adverse effect discrimination that accords with our intuition that

the two forms of discrimination are versions of the same phenomenon and amount to legal wrongs for at least some of the same reasons.

In the next two sections of the paper, I shall outline and defend an alternative view of adverse effect discrimination. I shall argue that it does involve fault or wrongdoing. But the fault has nothing to do with the agent's motives or purposes. Rather, it consists in interfering with another person's rightful claim to equal freedom. The core idea is this. Most of us enjoy certain freedoms of contract; and we believe, moreover, that everyone ought to have access to these same freedoms. These include both freedoms to enter into certain contracts and freedoms to deliberate about what we wish to contract for in a manner that is insulated from pressures stemming from certain extraneous facts about us, such as our race or religion. When someone else's actions interfere with your freedom, then the action can amount to adverse effect discrimination. Whether it does will depend on whether the measures that this person would have to take in order to accommodate you would constitute a greater intrusion into *his* freedom of contract. So the aim of prohibitions on adverse effect discrimination, on this view, is to protect each individual's right to a certain set of important freedoms of contract. It is to help secure the conditions of equal freedom. I shall argue that this is also the proper aim of legal prohibitions on intentional discrimination: intentional discrimination, too, is a legal wrong not because of the objectionable motives that accompany it, but because, like adverse effect discrimination, it amounts to an unjustifiable interference with another person's equal freedom.

I hope to flesh out this view in what follows, and to show that it provides a more coherent and attractive picture of discrimination than the ones we have so far considered. I shall argue that it resolves both of the apparent puzzles about adverse effect discrimination. And I shall try to show that it also answers to a further challenge, which I alluded to at the start of the paper. It explains why discrimination law does not constitute an unjustifiable interference with freedom of contract. It does not, on my view, because it interferes with one person's freedoms only in order to ensure that everyone has access to the same freedoms.

2.

Most of us enjoy certain freedoms in our contractual relations with others. Some are freedoms to *do* certain things. We are free to accept jobs that we have been offered, without having to turn them down because our religion prevents us from working on those days of the week that are required by the job; and we are free to take public transit without having to get off one stop early because the stop nearest our destination is not wheelchair accessible. Other freedoms that we normally enjoy are freedoms to have our *deliberations* about whom to contract with insulated from certain pressures, pressures caused by the way others view certain facts about us.¹³ We are free to consider which apartment to rent without having to factor in the contempt that the management might show for us because of our same-sex partner; free to decide where to shop without worrying about who will ask us to leave and breastfeed “somewhere more

¹³ See Seana Shiffrin, “Egalitarianism, Choice-Sensitivity and Accommodation”, in *Reason and Value*, ed. R. Jay Wallace et al. (Oxford: Clarendon Press, 2004) 270, for a detailed discussion of the moral importance of such deliberative freedoms.

discrete”; and free not to have to think at all about which sex’s washroom to use, because we need not fear retaliation from people who see cross-dressing or sex-change surgery as evidence of vice.

Most of us enjoy these freedoms because, when people enter into the provision of goods and services, or employment, or accommodation, they tend to arrange their affairs in such a way as to accommodate most people’s needs. The Sears Company schedules its workers’ shifts in such a way as to accommodate those whose Sabbath falls on a Sunday; the transit system designs its stations in such a way as to accommodate those who walk; and washrooms are assigned to one or the other of two sexes in accommodation of conventional views about gender. In other words, most of us enjoy these freedoms only because of constant accommodations, accommodations that are simply taken for granted and so remain invisible. But simply because these freedoms are in fact enjoyed by most of us, and the acts of accommodation that give rise to them, invisible, does not mean that there is no right to reasonable accommodation in play in these common cases. On the contrary, of the freedoms that I listed above, we believe not just that many people happen in fact to have these freedoms, but that each of us has a *right* to be accommodated in such a way as to have access to them ourselves.

We can see this by contrasting these freedoms of contract with others, which we do not think each person has a right to. There are many freedoms of contract that we do not think people can rightfully demand: for instance, the freedom to take up a job regardless of your qualifications for it, or the freedom to decide which apartment to rent only on the basis of the

aesthetic qualities of the apartment, without having to factor in whether you can afford it or whether it is close to your children's schools. Such pressures of finance and family needs are ones that we treat, to borrow a phrase from Lord Keith in the torts case of *Baker v. Willoughby*, as "the ordinary vicissitudes of life".¹⁴ When he wrote this, Lord Keith was not using 'ordinary' to refer to the frequency of these obstacles. That is, he was not making an empirical claim, the claim that because they are so commonplace, each of us must simply live with them. Rather, he was making a normative claim, the claim that there are certain burdens that we believe *ought* to lie where they fall, certain freedoms that people do not have a right to. The same is true, I am suggesting, of the freedoms we enjoy with respect to our contractual relations with others. There are some that we think are not freedoms that anyone can rightfully demand; though some of us may be privileged enough to enjoy them as a matter of fact. But there are other freedoms of contract that most of us do enjoy, and that we believe each of us has an equal right to enjoy.

I want to suggest that laws prohibiting adverse effect discrimination are one way of giving substance to these equal rights to certain freedoms in our contractual relations with others. For one of the things that adverse effect discrimination does is to interfere with a person's freedom of contract –either by preventing him from taking a certain action or by interfering with his freedom of deliberation by introducing extraneous pressures into his decision-making. Most often it interferes in both of these ways. Laws prohibiting adverse effect discrimination work to combat such interference. But significantly, they do not forbid every interference with another's freedom. All that they forbid are those interferences that satisfy two conditions: (i) they would

¹⁴ *Baker v. Willoughby* [1970] A.C. 467

not have affected the victim in this way had it not been for her possession of a trait that constitutes a prohibited “ground” of discrimination, and (ii) they cannot be eliminated by some reasonable accommodation of the victim, without causing undue hardship to the agent. We can see both of these conditions as ways of protecting equal freedom of contract for all.

Consider first the function of prohibited grounds of discrimination. Grounds of discrimination seem to express our judgments about which sorts of freedoms of contract we can rightfully claim of others, and which we cannot. For instance, every country that has anti-discrimination laws recognises race, religion, and sex as prohibited grounds, and this reflects our shared belief in the importance of being free to contract without obstacles or pressures based on these factors or on other people’s opinions about them. Some countries also recognise other grounds (in all or only some contexts), such as sexual orientation, disability, and citizenship. The types of traits that we treat as prohibited grounds differ widely. Some involve particularly important commitments, such as religion. Others, such as sex, are traits over which we have no (or no immediate) control; indeed, it may seem important to protect us from the impact of others’ views about these traits precisely because there is little we can do to alter them. Still other grounds, such as disability, involve some traits that can become the basis for important commitments, and others that we are simply saddled with and cannot alter: some disabilities, for instance, are developed into cherished aspects of one’s identity (think of those who are culturally deaf), whereas others are regarded as unfortunate obstacles (such as ADHD). But what matters for the purposes of discrimination law is not whether the trait is chosen or unchosen, an important part of our life or something we would rather be rid of. It is that, whatever the nature

of the trait, we believe that the conflicts between the needs of those who possess this trait and the needs of others, and the pressures stemming from others' beliefs about this trait, are not ones whose consequences the bearers of this trait can rightly be asked to bear. Their freedom of contract cannot rightly be curtailed simply because they possess these traits.

Grounds of discrimination thus reveal what we as a society take to be important freedoms at any given time. Our idea of this is of course evolving, and has developed considerably over the years. Many countries now believe, for instance, that "gender identity" should be added to the list of prohibited grounds of discrimination. We used to assume that transgendered persons had no right to question the institutional structures and assumptions that excluded them: we found their desires immoral and their behaviour, offensive, and we believed that for this reason, they deserved the disadvantages they faced. But we now think that these assumptions are false: transgendered persons deserve freedom from social obstacles and pressures that arise from the lack of correspondence between their biological sex and the gender they identify with or from their attempts to achieve such correspondence.

If this is a correct analysis of the function of grounds, then it helps to explain why so many human rights tribunals and courts have faltered in attempting to locate a single criterion for something counting as a prohibited ground. On the view I have been suggesting, whether some trait should be recognised as a prohibited ground is a normative question whose answer depends on whether people have a right to be free from the sorts of social pressures and obstacles that are encountered by those with that trait. Because there is no reason to suppose that we value all

freedoms for the same reasons, there is likewise no reason to think that all grounds of discrimination will meet some single criterion for constituting a ground. On the contrary, it seems quite likely that the reasons will be diverse. We think everyone ought to be free to make contracts and deliberate about them without facing pressures based on the exigencies of their religion, and one reason for this may be that religious commitments are such an important part of many people's lives that we believe that they should be able to live in light of these core beliefs, without facing undue burdens because of them. But the same explanation will not do for the ground of "colour". For many people, their skin colour is a salient fact about themselves only to the extent that it is a means by which others have belittled them and others whom they have grouped together with them on this basis. It does not, in these cases, represent anything chosen by them or anything they identify with; so it does not always tie together a community of people who share anything more than a history of stigmatization and stereotyping. Yet we still think that people deserve protection from exclusions on the basis of skin colour.

I have argued that prohibited grounds show us which freedoms we feel individuals have a right to claim in their dealings with others. I want now to turn to condition (ii), the requirement that we must make reasonable accommodation of others, but only up to the point of undue hardship. On the view that I am proposing, the function of this requirement is to ensure that the freedoms of those who face discrimination are protected, but not at the expense of even greater intrusions into the freedoms of the discriminators. In other words, the requirement ensures that potential victims of discrimination are given access to the same freedoms that others commonly have, but only insofar as this is compatible with the agent's retaining his access to such

freedoms. If you are a full-time employee in my bookstore and I have a policy of requiring all full-time employees to work one Saturday a month, I must make an exception for you if you are a Seventh Day Adventist and need Saturdays for religious observance –at least, as long as my store is large enough that other employees can cover for you. For requiring me to accommodate you gives you the important freedom to work at a job that you are qualified for, and the freedom to deliberate about the job without having to consider the impact of your religion on your work (or the impact of your work on your religion). And the only freedom of mine that this interferes with is a freedom which we do not think everyone has a right to, namely the freedom to dictate every aspect of the running of my business myself. However, if you are my only employee and my bookstore would be forced to close on Saturdays if you did not work that day, and if such closure would substantially interfere with the viability of the bookstore, then I can reasonably demand that you work on Saturdays. The law, then, requires reasonable accommodation, but only to the point of undue hardship, because it is concerned not just with the freedoms of those who face discrimination but also with the freedoms of those whose actions may result in it. It is concerned to ensure that each of us is given no more and no less freedom of contract than the majority of other people enjoy.

I have tried to explain how we might see prohibitions on adverse effect discrimination as attempts to promote equal freedom of contract for all. But one might object that I have not yet made clear how this picture of adverse effect discrimination makes sense of it as involving wrongdoing, in such a way as to support both our intuition that the agent can legitimately be coerced into shouldering the costs of reasonable accommodation, and our intuition that the

victims of discrimination have been personally wronged by the agent. This objection might take one of several forms. I shall discuss these, in turn.

First, one might be concerned that as I have described adverse effect discrimination, there does not seem to be any intent requirement. All that seems to matter is that one's behaviour has the effect of interfering with another person's freedom to a greater extent than any accommodations you might make for her would interfere with your freedom. But if one's behaviour need not be intentional in order to amount to discrimination, how can it then amount to wrongdoing (or, more particularly, to a wrong that one does to someone)?

This form of the objection seems misconceived. For adverse effect discrimination is intentional in the only sense that might be relevant to legal wrongdoing: the discriminator always intends the action or rule that gives rise to discrimination under *some* description, even if he does not initially intend it under the description of "excluding Sikh officers from ceremonies because of their religion", or under whatever description makes it apparent that it is a case of adverse effect discrimination. The practices or rules that result in adverse effect discrimination are never mere pieces of behaviour. They are always intentional under some description; or, as we might otherwise put it, they are always *actions*. It is true, of course, that they are not generally chosen under the description that makes them appear discriminatory. But, as long as they are intentional under some description, they can still be regarded as the perpetrator's actions.

However, one might now interject –and this is the second form that the objection might take—that this is not sufficient to show that adverse effect discrimination involves wrongdoing. What seems problematic, one might suggest, isn't that the action is not intentional in the right sort of way; it is that the objectionable features of the action, on my view, seem to consist solely in its effects, and not at all in the nature of the action leading to these effects. Adverse effect discrimination, on the view that I have articulated, is objectionable because it results in the denial of equal freedoms to some people. But this seems to be just a matter of the effects of an action, and to have little or nothing to do with what the agent was *doing*.

In responding to this version of the objection, it may help to compare adverse effect discrimination to a tort that shares much the same form: the tort of defamation. This tort involves publishing a statement that has the effect of damaging another person's reputation, or image in the community. Like adverse effect discrimination, defamation is a legal wrong because of its effects, rather than because of the kind of action that the agent performed. Hence, the plaintiff need not show that the defendant published the statement maliciously (nor even that he published it intentionally; though unintentional publications can only be defamatory if the publication was done negligently, because the defendant could not otherwise legitimately be regarded as their author). Nor does the wrong of defamation seem to lie in any facts about the *way* in which the statement was published. Rather, it lies simply in the harmful effects on the individual's reputation. As one Canadian commentator has said, defamation is “defined by the

character of the injury it causes, rather than by the character of the insult which it produces”.¹⁵ 27

The same could be said about adverse effect discrimination.

But is this not problematic? In the case of defamation law, it might be. For defamation law contains no other components that might help us locate wrongdoing on the part of the defendant. Although there are a number of defences to defamation –such as the truth of the statement published; or the “privilege” of publishing defamatory statements in the discharge of certain private or public duties; or the need to allow for “fair comment” on matters of public importance—these defences do not focus on legitimating features of the defendant’s *actions*. So they do not work to suggest that, in the absence of these legitimating features, we can impute wrongdoing to the defendant. Rather, the defences seem to articulate a series of policy-based exceptions that focus on the importance of free speech and its ability, in certain circumstances, to outweigh the harmful effects of the defamatory statement on the plaintiff’s reputation.

However, in the case of adverse effect discrimination, there is a further feature that helps us explain how the agent’s actions could amount to wrongdoing. This is the requirement that the agent must provide only *reasonable* accommodations, and only up to the point of *undue hardship*. Because this requirement allows us to deny that discrimination has occurred at all in cases where there is no way of reasonably accommodating others short of undue hardship, it implies that in all cases where there *is* discrimination there must have been reasonable measures of accommodation available to the agent and he must have failed to take them. Hence, it is not

¹⁵ R. Browne, *The Law of Defamation in Canada*, 2nd ed. Carswell, 1999; my italics.

really accurate to say that adverse effect discrimination consists solely in the creation of certain effects. A better suggestion is that it consists in the agent's failing to make reasonable accommodations for others. It follows that those who engage in adverse effect discrimination are acting *unreasonably*. We can therefore impute unreasonableness to agents who persist in discriminatory actions. Their fault, then, can be understood as readily as the fault in any negligence-related tort.

Or can it? One might object –and this is the third version of the fault objection—that although we can readily understand why most of the actions that amount to negligence in the law of torts are unreasonable, and hence involve fault, the same is not true of adverse effect discrimination. Why should we regard the failure to accommodate those who are excluded by my actions on the basis of certain character traits as unreasonableness on my part, as something that is my problem and falls on my shoulders to rectify? Another way of putting this third version of the fault objection is to say that it is unclear how my own account of adverse effect discrimination does not suffer from the same defects that, earlier, I found in Gardner's view. On Gardner's view, those who engage in adverse effect discrimination are simply poor agents of distributive justice: all that they have failed at is appropriately discharging the state's obligation to redistribute resources and opportunities in accordance with the preferred principles of distributive justice. I argued earlier that it is difficult to see how this could be their fault or amount to wrongdoing in more than the positivist sense that it is actually treated as a legal wrong. One might object that the same can be said of my view. Why is it unreasonable, and hence my fault, if I choose not to accommodate you? Surely all that I am refusing to do in such

cases is come to your aid or benefit you, by making available to you the freedoms that most people possess. But why is that my responsibility?

There is an important difference between Gardner's view and mine, which I think can help to answer this objection. On Gardner's view, adverse effect discrimination is objectionable because it does not benefit certain people in circumstances where the state owes them a benefit. Adverse effect discrimination amounts, on Gardner's view, to a failure to redistribute resources and opportunities in the optimal way: hence, on this view, it is objectionable simply because of the failure to benefit others. On my view, however, the objection to adverse effect discrimination is not that it fails to benefit them. It is that it interferes with their freedom; and this freedom is something to which they have an equal right. Each of us has a claim on others to have the same important freedoms of contract that the majority of us enjoy. And so I wrong you if my rules or actions prevent you from having one of these freedoms. It is of course true that, when I interfere with your freedom, I normally harm you or at least fail to benefit you. But what makes my action into a legal wrong –what enables us to see it as my fault, from a legal standpoint—is not the mere fact that I have harmed you or failed to aid you, but rather the fact that I have interfered with a freedom that is rightfully yours. On my view, then, the failure to accommodate others is unreasonable, and hence amounts to wrongdoing, because it is an interference with their rightful freedom. And for the same reason, it is a personal wrong against the victims of discrimination. It is not simply an instance of misconduct that has the effect of harming them: rather, in interfering with someone else's rightful freedom, I commit a personal wrong against them.

I have tried to show that by construing adverse effect discrimination as an interference with our equal right to certain important freedoms, we can make sense of it as involving wrongdoing –and moreover, the kind of wrongdoing that involves a personal wrong toward the victims of discrimination. We can therefore explain why it is that the victim of adverse effect discrimination can rightly claim that she has been wronged by the discriminator, and why it is just to require the discriminator to cover the costs of reasonable accommodation.

I now want to suggest that this is a plausible way of understanding why intentional discrimination is a legal wrong, as well --and that, as a result, my account of adverse effect discrimination helps us to make sense of these two kinds of discrimination as two forms of the same phenomenon. Recall that intentional discrimination arises when actions or rules are adopted with the purpose of excluding some people because they possess a certain trait, and this trait is a prohibited ground. Hence, intentional discrimination differs from adverse effect discrimination in that its purpose is to exclude, whereas the exclusionary effects of adverse effect discrimination are not part of its purpose. On the common view of intentional discrimination that we examined earlier, what makes this form of discrimination into a legal wrong is the presence of one of two further features: either the agent's sole motive is malicious, as she excludes others in order to harm them, or it is prejudicial, involving a belief in the inherent inferiority of these people. This view is problematic in at least two respects. On the one hand, some exclusions are overt and purposive and seem objectionable, but they are not in any way motivated by malice or prejudice. Consider, as an example, a case that recently arose in

Vancouver.¹⁶ It concerned a rape crisis centre that refused to allow someone to volunteer as a peer counsellor at the centre because she was a male-to-female transsexual. Because the crisis centre explicitly excluded her from volunteer work based on her gender identity, the case is one of deliberate exclusion on the basis of a trait that is a prohibited ground of discrimination. However, the managers of the crisis centre were not motivated by any desire to cause harm to this person; nor did they think that her being a transsexual somehow lowered her value as a human being. Their main reason was that, having not been born or raised as a woman, she lacked substantial personal experience of the particular kind of discrimination that had been faced by the women she would counsel, and so she would in their view be a less effective counsellor; and they argued, in addition, that she still retained enough of the appearance of a man that battered women might find it traumatizing to have to discuss their experiences with her. This reasoning may seem both unimaginative and insensitive; but it surely does not involve malice or prejudice.

The second, and more important problem with the common view of intentional discrimination is that, even where malice and prejudice are present, it is unclear that they can play the type of role that the common view envisages them playing. According to the common view, it is these motives that turn intentional discrimination into a legal wrong. But, at least in tort law, the presence of certain motives is rarely enough to turn an otherwise lawful action into a legal wrong – unless, that is, the motive changes the character of what is done and enables us to describe the action as a different action, an action that is wrongful. So, for instance, in the case

¹⁶ *Vancouver Rape Relief Society v. Nixon*, 2005 B.C.C.A. 601 (British Columbia Court of Appeal).

of intentional torts, certain intentions play a role in rendering an action wrongful because they make a difference to the kind of action it is. If I inadvertently hit you with my backpack while removing it, I have not assaulted you. But if I deliberately swat you with it while removing it, I have. In this case, my intention to hit you makes a difference to the legal status of my action because it turns what would otherwise have been an unintentional touching into an action of a different kind: an unauthorized intentional touching, or an assault. And an unauthorized intentional touching amounts to a legal wrong because it is a way of using you; whereas, when I bump you inadvertently, I don't use you at all. But aside from cases in which the presence of a certain intention changes the nature of the action performed, mental states are rarely treated in tort law as sufficient to render an otherwise lawful action unlawful. Most courts, for instance, deny that malice is ever, on its own, sufficient to convert lawful actions into legal wrongs.¹⁷ So the common view of intentional discrimination sits uncomfortably with the normal role of motives in tort law.

If we model an account of intentional discrimination on my explanation of what makes adverse effect discrimination into a legal wrong, we avoid both of these difficulties. We can say that intentional discrimination differs from adverse effect discrimination simply in the fact that its purpose is to exclude, whereas the exclusionary effects of adverse effect discrimination are not part of its purpose. Intentional discrimination may or may not, on this view, involve

¹⁷ As Denise Reaume has argued, it is consistent with this point to accept that malice can defeat a privilege in certain contexts. For in such contexts, it is features of the action other than the agent's malice that make that action into a legal wrong. The role of malice here is not to turn the action into a legal wrong, but rather to show that the agent was not in fact acting in accordance with the rationale for the privilege, and hence, that the privilege should not operate as a defence. See Reaume, "Harm and Fault in Discrimination Law: The Transition from Intentional to Adverse Effect Discrimination", (2001) 2.1 *Theoretical Inquiries in Law* 349 at 360, note 19.

objectionable motives. But what makes it into wrongdoing at law is never these motives. It is, rather, wrongful in the same circumstances, and for the same reasons, that adverse effect discrimination is wrongful: namely, because and to the extent that it constitutes an objectionable interference with another person's rightful freedom of contract. On this view, of course, the distinction between the two forms of discrimination is not of very much legal importance. It may be of minor importance: victims of intentional discrimination may suffer more when they are purposively excluded, because such exclusions may constitute more of an insult to them (in ways that I shall discuss further in section 3 of the paper). So it may be that special damages should be awarded in such cases, to reflect one or another of these additional harms. But these harms should be regarded as further effects of discrimination for which the agent becomes responsible because he is already responsible for the discrimination, in the same way that tortious agents become responsible for the harms resulting from their wrongful actions. They should not, I am suggesting, be treated as part of the reason why the discriminatory action was wrongful.

Before I go on to consider some objections to my view, I want to note one significant implication of the view. It concerns a worry that I alluded to early in the paper. This is that discrimination law constitutes an unjustifiable interference with freedom of contract and association. Richard Epstein has eloquently articulated this worry in a number of places; and he laments the fact that our society has become "a world in which the antidiscrimination norm is

regarded as superior to the principle of freedom of association”.¹⁸ My account suggests that this way of thinking about the relationship between discrimination law and freedom of association is mistaken. Discrimination law does not simply function as a constraint upon freedom of contract and freedom of association. Rather, it constrains some people’s freedom of contract and association in order that others, too, can share in these same freedoms to an equal extent. Moreover, if I am right that most of us have these freedoms only because others are constantly, and tacitly, making certain accommodations for us –if I am right, that is, that the basic case is actually one of accommodation, and that those of us who do not have to launch discrimination complaints have only avoided the need to do so because our rights have *already* been respected—then it only looks as though the burdens of discrimination law fall unequally on those who end up having to accommodate persons with traits marked out by prohibited grounds. Discrimination law only asks these individuals to do what we *all* implicitly do most of the time: that is, make reasonable accommodation for others, so that they too may have access to important freedoms of contract.

Of course, this is not a response that would convince Epstein. On his view, freedom of contract is important only instrumentally, insofar as it helps to maximize the overall level of utility in society. Because Epstein values freedom only instrumentally, and only insofar as it helps maximize overall utility, he is not at all concerned with the way in which freedoms are distributed across a society. But for those of us who take a different view of the value of

¹⁸ “Castle and the Civil Rights Laws: From Jim Crow to Same-Sex Marriages” (1994) 92 Mich L. Rev. 2456 at 2470. See also *Forbidden Grounds: The Case Against Employment Discrimination Laws* (1992), and “Standing Firm on Forbidden Grounds”, (1994) 3.1 San Diego L. Rev. 1.

freedom of contract and the purpose of the state's attempts to protect it, my account may help to quell worries about the extent to which discrimination law interferes with it.

3.

One might at this point object that, even if my account succeeds in explaining how adverse effect discrimination could involve wrongdoing, it overlooks the most important feature of any discriminatory action. This is the message that such actions send to others about the value of those who have been excluded: the contempt or disregard that the actions express. Intentional discrimination—at least, where it involves malice or prejudice—obviously expresses such contempt. But even adverse effect discrimination may send a certain message about those who are excluded. To fail to accommodate others in circumstances where you could do so without undue hardship may, in some circumstances, send the message that their freedoms are not worth worrying about; and so your actions may express the view (even if you do not hold this view or intend to convey it) that these other people are second-class citizens. Some legal scholars have recently argued that it is really the fact that discriminatory actions constitute such objectionable expressions that explains why discrimination should be legally prohibited.¹⁹ According to these expressivists, when we enter the public sphere by offering employment or accommodation or services to the public, we must act in a way that expresses equal concern and respect for others.

¹⁹ Anderson and Pildes, “Expressive Theories of Law: A General Restatement”, (2000) 148 U. Pa. L. Rev. 1503; Hellman, “The Expressive Dimension of Equal Protection” (2000) 85 Minn. L. Rev. 1.

Discriminatory actions fail to express this. Hence, a natural way of understanding why they are legally prohibited is in terms of their failure to express the right sorts of attitudes about others.

Indeed, an expressivist might argue that we can only properly understand the function of *grounds* of discrimination if we adopt this explanation of why discrimination is legally objectionable. Grounds of discrimination are not, the expressivist might say, just any trait of which we think that the person bearing the trait should not have to shoulder the costs associated with it and others' attitudes about it. On the contrary, grounds are generally traits that have been used to express demeaning messages about others: think of race, sex, religion, or indeed any of the traits that has made it onto any country's list of prohibited grounds of discrimination. This is why, when the Canadian Supreme Court was confronted with the claim of a self-styled "marihuana freedom-fighter" that a "taste for the recreational use of marihuana smoked through a glass water bong" should be recognised as a ground of discrimination, the Court expressed a mixture of incredulity and outrage, saying that this would be to "create a parody of a noble purpose" and that it would "trivialize this list to say that 'pot' smoking is analogous to gender and religion . . .".²⁰ Even if marihuana freedom-fighters face the same level of exclusion as individuals who possess traits that have been recognised as prohibited grounds, they do not suffer from the same stigmatization and insult.

²⁰ *R. v. Malmö-Levine* [2003] 3 S.C.R. 571 at para. 185.

We might doubt, I think, whether the stigma associated with being a marijuana freedom fighter is really any less than the stigma associated with exclusions based on sex or religion. Indeed, part of the reason why this example seems humorous is that most of us do think of the marijuana “freedom-fighter” in terms of rather demeaning stereotypes: he is likely a high-school drop-out, is obsessed with a trivial activity, and alternates between flamboyantly crying out for freedom and sinking into a drug-induced stupor. In other words, we do stigmatize him, but we think it is acceptable to do so because we think he deserves it. So perhaps the difference between “a taste for marijuana” and recognised grounds of discrimination is not that actions that exclude people on the basis of the former trait express any less disdain; it is that this disdain does not bother us because we think it is deserved. Of course, the expressivist might rejoin that, to the extent that the disdain is deserved, the action does not express a lack of equal concern and respect. But this response commits the expressivist to a rather more complicated account of when actions fail to express equal concern and respect than expressivists sometimes give. It is not enough, on this view, for the actions simply to convey a disdainful or contemptuous message. It must also be the case that this message is not deserved. This is not necessarily a problem for expressivists; but it does suggest that the account needs to be more complicated than are some of the expressivist accounts offered.

Before we turn to a more detailed discussion of expressivism and the reasons why, in my view, it seems problematic, it is worth noting that my account of discrimination, too, can explain both why the traits that are recognised as grounds seem to be associated with stigmatization, and why, as the case of “a taste for marijuana” suggests, not all traits associated with stigmatization

are grounds. On my view, the reason why “a taste for marihuana” seems not to fit as a ground of discrimination is that, even though it is lamentable that marihuana-users are often stigmatized, we do not think that the freedom to smoke marihuana recreationally regardless of the consequences is a sufficiently important one to curtail others’ freedoms in order to secure it. When a particular freedom of contract *is* sufficiently important to warrant protection through discrimination law, however, its importance to us, combined with the fact that most people enjoy this freedom, means that the denial of this freedom to some people may take on a certain social significance: it may suggest or imply that these people are second-class citizens. Hence, my view, too, can explain the stigmatization associated with discrimination. It just does not treat this as the reason for legally prohibiting it.

I want now to suggest that expressivist views of why discrimination is legally prohibited are in a number of respects problematic. My arguments are not intended to cast doubt on expressivism as a general theory of value or as an explanation of why discrimination is objectionable from a moral standpoint; and they may still leave room for expressivism as a legal theory that applies to other areas of law. I shall be arguing only that, as an explanation of why discrimination should be legally prohibited, expressivism is inadequate.

It will help to begin by clarifying several features of expressivist views of discrimination. Expressivism, we have seen, focuses on the attitudes that actions convey. As Anderson and Pildes have emphasized, an action can convey a certain attitude without being caused by it. It can also convey that attitude even if the agent does not intend to communicate it to others. So

the idea of what an action expresses is not to be confused with the idea of what intentions cause it, or what the agent intends to communicate through it. Indeed, this is why adverse effect discrimination can amount to an expression of a lack of equal concern and respect for others, even though it is not usually intended as such, and even though it is usually no part of the agent's purpose to communicate lack of regard for others. But if it is not the agent's intention that tells us what his action expresses, what does? The answer that most expressivists give is that what an action expresses depends on the social context and our shared sense of what different actions mean in that context. There are, as Anderson and Pildes state, "objective criteria for determining the meanings of action". These criteria are what we look to, when we determine what an action expresses. Hence, expressivism "holds people accountable for the public meanings of their actions".²¹

A natural question that this account raises is: why should we hold people accountable for these public meanings of their actions? How is it that they can be justifiably *coerced*, simply on the basis of what their actions publicly express –when this public expression is not a function of the agent's own intentions, and not even something that a reasonable person in the agent's position would always recognise?

There are several problems here. One grows out of the objective conception of expression that expressivists employ. Anderson and Pildes suggest, I have said, that the objective meaning of an action is a function, not just of what a reasonable person in my position

²¹ Anderson and Pildes, p.5.

would take it to mean, but of what it actually *does mean* given the relevant social conventions and the surrounding context. But it follows that it is possible for an action to express something that the agent cannot have been aware of, and that even a reasonable person in his position would not have been aware of. And this leaves it mysterious how we can regard these expressions as being the agent's fault. How can something be my fault, if I did not know that I was expressing it, and a reasonable person in my position would not have known? It is unclear, then, that an expressivist account of discrimination could make sense of discrimination as involving fault or wrongdoing.

Another problem posed by the suggestion that we can justifiably be coerced on the basis of what our actions express is that expressions are not, at least in the law of torts, the sort of thing that generally attracts legal prohibitions. Recall that, earlier in the paper, I noted that tort law does not generally treat mental as the object of its prohibitions. A person's intentions can of course make a difference to the kind of action he performs; but when they do this, it is still the action, and not the intention, that tort law prohibits. The same sort of point can, I think, be made about the attitudes that an action expresses. Expressivism begins with the plausible idea that many actions are expressive of certain attitudes.²² But it does not follow that all of the norms that regulate action must be norms that regulate the expressive meanings of actions. Some or all legal norms, for instance, may simply be focussed on ensuring that certain consequences are not

²²Note, however, that Anderson and Pildes make a stronger claim, and one which seems to me dubious: they claim that *all* actions are expressions. But if I am startled into pulling the trigger on a gun, does my action express anything?

produced --in the case of discrimination law, consequences that involve interferences with other people's equal freedom.

Perhaps most importantly, expressivism sits uncomfortably with certain fundamental features of discrimination law. Firstly, the same content that is expressed by certain actions which the law does prohibit as discrimination can equally well be expressed through other actions, actions that are not regulated by discrimination law and are not subject to any legal sanctions. Suppose I own a small pharmacy. I am not free to deny jobs at my pharmacy to women simply on the grounds of their gender: this would constitute discrimination, and would be prohibited by law. But I am perfectly free to stage a peaceful demonstration against hiring women pharmacists on the street outside and to wear a placard stating that women should not be hired as pharmacists. This would not be prohibited by private sector discrimination law, since it is not something I do in the context of being an employer. And this suggests that what is wrong with my denying jobs to women as an employer is not reducible to what my action expresses. For surely most of this is also expressed by my perfectly legal placard and my demonstration. Secondly, it seems significant that although our constitutional entitlement to freedom of expression protects many ways of expressing stigmatizing or exclusionary claims about others (such as my marching up and down with my placard), no role at all is given in discrimination law to the discriminator's interest in freedom of expression. At no stage of the proceedings is the alleged discriminator allowed to put forward the claim that his interest in freedom of expression would be harmed if he were required to accommodate the victim; and there is no opportunity for the adjudicating body to weigh this interest in freedom of expression against the victim's interest

in being accommodated. This suggests that the focus of discrimination law is not on the discriminator's action *qua* expression.

I have tried to suggest that expressivism runs into difficulties as a theory of why discrimination can justifiably be prohibited by law. But I want now to suggest that expressivism does point us in the direction of an important truth about many cases of intentional discrimination, and about the ways in which the presence of malice or prejudice can aggravate the wrong that is done by discriminatory actions. I noted earlier that expressivism starts from the plausible claim that many actions constitute expressions –and in particular, that many discriminatory actions constitute expression of contempt for others. This seems most plausible as a description of intentional discrimination, or more exactly, those kinds of intentional discrimination that are motivated by malice or prejudice. In such cases, the agent has not simply overlooked a particular person and her needs, as he may have done in the case of adverse effect discrimination; rather, he has actively ignored or trampled on them, and this does express contempt and disdain for her. This contempt and disdain clearly cause additional anguish to the victims of discrimination; and they may also increase the social stigma associated with possessing the trait in question. So in such cases, the contemptuous nature of the action *qua* expression seems to aggravate the wrong that is done. Not only does the discrimination interfere with the victim's freedom, but it causes her additional emotional suffering and results in additional social stigmatization of her. This may be why many jurisdictions award special “dignitary” damages in cases of intentional discrimination, to compensate the victim for these additional losses. If I am right about this, then expressivism helps us to see that when

discrimination is done intentionally, and when it is motivated by malice or prejudice, the contempt that it expresses aggravates the wrong that is done to the victims. It is important to note, however, that it is simply the particular *way* in which the action is done in such cases that makes it an expression of contempt, and that this is not an essential feature of the action as a discriminatory action. So the fact that a discriminatory action that is motivated by malice or prejudice expresses contempt for the victims may aggravate the wrong or add a further consequential loss which we think the discriminator can rightly be held accountable for. But it does not constitute the primary wrong that is done in cases of intentional discrimination.

There remains the question of whether adverse effect discrimination, too, can ever be aggravated by the way in which it is done and the message that it thereby expresses. I suggested earlier in this section that, when we enter the public spheres of employment or provision of goods and services but fail to consider the impact of a rule or policy on a particular group of people, this can sometimes send the message that their needs are not worth considering. And it is easy to imagine this causing emotional suffering and social stigmatization quite similar to – though likely not as severe as—that caused by discrimination motivated by prejudice or malice. If this is right, then it may be that the special dignitary damages that we now reserve for cases of intentional discrimination should be open also to victims of adverse effect discrimination, though one would expect lesser amounts to be awarded in such cases than in those cases of intentional discrimination that involve malice or prejudice. This is a question that requires further consideration, and a more detailed analysis of the message expressed by adverse effect discrimination than I can give here.

In this paper, I have argued that we should understand both intentional and adverse effect discrimination in terms of wrongful interference with another person's equal freedom. If we do so, we can make sense of discrimination as a single, coherent phenomenon. And we can explain why it is that those who discriminate seem to wrong those whom their actions exclude, and why they can justifiably be coerced into covering the costs of reasonably accommodating them. This account also provides a clear explanation of why discrimination law is not an unjustifiable interference with freedom of contract: it is not, because it limits some people's freedoms only in order to enable everyone to enjoy the same freedoms. As the last section of the paper suggests, there remains work to be done in clarifying the precise role that should be given in discrimination law to the emotional suffering and stigma resulting from the way in which the discriminatory act was performed. I have argued that these factors aggravate the wrong rather than constituting it, and I have suggested that although they are most clearly present in those cases of intentional discrimination that involve malice or prejudice, they may also be present in other cases of intentional discrimination and in cases of adverse effect discrimination as well –though I am not certain that this is correct. But I hope that this paper has provided at least the beginnings of a workable account of discrimination and has shown why it is a perplexing and philosophically interesting phenomenon.