

Are bad things worse when they happen to you because of who you are?

Discrimination and Status Wrongs

Robin Wisse

MPhil Stud Philosophy

University College London

Supervised by Dr Robert Simpson

Abstract

Intuitively, discrimination and identity-based harms are distinctively wrong. Why would that be so? I argue that we have no reason to think that socially incurred harms are distinctively wrong apart from our intuition that it is so (chapter 1). I then move onto status-based accounts of the wrong of discrimination and identity-based harms as presented by Chambers (2008), Moreau (2020) and Hellman (2011). I argue that while they are right to focus on status, they fail to capture the distinctive and personal wrong incurred by people in these circumstances for two reasons. First, I argue that they mistakenly rely on risk of harm, and put forward an alternative status wrong account (chapter 2). Second, I argue that Hellman's account rightly relies on interpretative value judgements, but that it lacks a convincing account of how these can be objectively valid and therefore can be taken seriously as a source of moral knowledge (chapter 3).

I trace both of these shortcomings back to an insufficient metaethics on which objectivity and subjectivity are mutually exclusive. If we want to maintain and develop the view that identity-based mistreatment is distinctively wrong we need a metaethics to underwrite it. To this end, I put forward Crary's alternative Wittgensteinian metaethics (Crary, 2016). This shift in metaethics allows me to rephrase the question. Instead of asking '*Why* are identity-based harms and wrongs distinctively and personally wrong?', I can now ask: 'Does this intuition tell us something about how the world really is, or is it just a misleading sentimental response?' I argue that there is no reason to doubt that this intuition tells us what the world is really like by drawing on and developing Crary's metaethics (chapter 4). I develop Crary's view by showing how it can reject racist literature. Finally, I show how the view can help us assess the meaning of social practices like affirmative action (chapter 5).

Impact Statement

This thesis aims to improve on status-based accounts of the wrong of discrimination, and provide a metaethical foundation that can underwrite them. As such it engages with and connects two academic fields: work on discrimination and identity-based harms, and work on what I call Wittgensteinian metaethics.

The thesis aims to benefit the academic discussion in three main ways. First, it identifies shortcomings of existing status-based accounts of discrimination and identity-based harms, as put forward by Chambers (2008) Moreau (2019, 2020) and Hellman (2011, 2018). Most notably, it argues that we cannot capture the wrong of discrimination in terms of (a risk of) harm, but that instead, we should be looking to develop a status wrong account. Second, it draws attention to the lack of adequate metaethical foundations of these views and the need for such a foundation. Third, it puts forward and develops a Wittgensteinian metaethics that can underwrite such accounts based on Crary (2016). In doing so it recommends a shift in research methods: instead of looking for facts that make discrimination wrong, we should investigate whether the intuition that discrimination is wrong tells us what the world is really like, making use of our moral imagination and affective responses. This change in metaethics and method enables philosophers to take moral intuitions and the subjective and emotional experiences seriously, instead of treating them as inherently suspicious.

Discrimination and social inequality more broadly are public issues, and the ideas discussed in this work are meant to have benefits beyond academia, specifically in public discourse, culture and law. I would like to highlight four such ideas. First, discrimination is distinctively wrong because it casts people as inferior. This means it is only distinctively wrong to discriminate against oppressed people, not privileged ones. We shouldn't be talking about 'gender' or 'race' discrimination, but of discrimination against women and Black or brown people. We should also expand our notion of discrimination to include oppressed identities, such as working class people and immigrants. Second, discriminatory treatment without a risk of harm is still morally wrong. Third, overly emotional or sentimental arguments aren't inherently suspicious, and qualify as serious, rational arguments. Fourth, affirmative action isn't necessarily demeaning. The dignity of a student depends not on whether they were admitted on an affirmative action programme, but on how they are treated by their university.

To achieve this, parts of this work might be published in academic journals, newspapers, magazines, blogs or Twitter, and presented at academic conferences or meetings open to the general public.

INTRODUCTION – ARE BAD THINGS WORSE WHEN THEY HAPPEN TO YOU BECAUSE OF WHO YOU ARE?	5
CHAPTER 1 – ARE IDENTITY–BASED HARMS DISTINCTIVELY WRONG?	8
1.1 Lottery–world	9
1.2 Harm and unequal social norms	10
1.3 The Social Condition	13
1.4 Responsibility	16
1.5 The Inequality Condition and Status Harm	20
1.6 Conclusion	22
CHAPTER 2 – STATUS HARMS AND STATUS WRONGS	24
2.1 Moreau’s Unfair Subordination	25
2.2 Status Harm and Social Meaning	28
2.3 Hellman’s Demeaning	29
2.4 Harms and Wrongs	34
2.5 Conclusion	38
CHAPTER 3 – SUBJECTIVITY AND OBJECTIVITY	39
3.1 Interpretative value judgements	40
3.2 Modest Objectivity	43
3.3 Wittgensteinian Metaethics	48
3.4 Conclusion	54
CHAPTER FOUR – GETTING IT RIGHT	56
4.1 Reading <i>The Camp of the Saints</i>	58
4.2 Authoritative Perspectives	63
4.3 A problem with attention	69
4.5 Conclusion	71
CHAPTER FIVE – AFFIRMATIVE ACTION	72
5.1 <i>On Beauty</i> : Patterns of attention	73
5.2 Conclusion	78
CONCLUSION – NO NEED FOR JUSTIFICATION	80
REFERENCES	81

Introduction – Are bad things worse when they happen to you because of who you are?¹

Our world is one of structural social inequality. Because of this, who you are matters: when some groups are structurally privileged and others structurally disadvantaged, your social identity can affect your life for the better or the worse. Here are some examples: Being a woman makes you vulnerable to sexual harassment and assault. Being poor makes you less likely to survive cancer. Being Black makes it more likely that doctors won't take your pain seriously. These social identities intersect. Black women are subject to special types of racist and sexist abuse. Poor women are even more vulnerable to domestic abuse. And the Black working class experience more economic anxiety.

In other words, social inequality is vast and many people have social identities in virtue of which they are worse off in unique ways. The fact that people structurally incur these wrongs and harms in virtue of their identity seems profoundly wrong. What's more, the people who find themselves in this situation seem to have incurred a distinctive, personal wrong. The wrong seems personal because these people seem to be able to claim that they are wronged as individual people. And the wrong seems distinctive because it seems to be an additional wrong on top of the different kinds of harm that might also be involved. This is the central intuition I want to explore in this thesis: why is it especially bad to be wronged or harmed in virtue of your social identity?

Here is one slightly more detailed, first-personal example of the intuition I have in mind:

Frozen out of a friendship: A male friend told me that I wasn't interesting or deep enough to have an intimate friendship with because of my gender. When I explained that I was hurt and offended, he rushed to ensure me that it was nothing personal, and that he wasn't being sexist. He had nothing against me or against women on principle, he told me, but it was just a fact that women were socialised to be less interesting, and thus not good enough for him. When I insisted he was treating me badly, he kept explaining that my feelings were misplaced. This 'discussion' lasted about half an hour. I still feel humiliated and angry. He made me feel that I had to prove my worth to him, and made me feel like I was going crazy.

Identity-based wrongs and harms come in a wide variety. First, there is direct discrimination: someone is treated worse based on their (perceived) social identity. Second, there is indirect discrimination: someone is treated worse based on a trait linked to their social identity. Third, there

¹ I am grateful to Rob Simpson, Seb Bishop, Filippa Ronquist and Euan Alison for their comments and encouragement, and to Stefan Kolling for his support.

are cases in which people harm themselves because of their social identity. And lastly, sometimes the causal links between someone's social identity and their disadvantage aren't clear. In those cases it's clear that people with a particular identity are disadvantaged, but unclear precisely which mechanism causes their disadvantage. When I say that it's especially bad to be wronged in virtue of your social identity, that claim is supposed to encompass this variety of cases.

I begin by engaging with Clare Chambers' work on identity-based harms (chapter one) and status-based theories of the wrong of discrimination by Sophia Moreau and Deborah Hellman (chapter two). My discussion of these authors is guided by three questions:

Q1: What is special about identity-based social inequality, compared to inequality based on other grounds?

Q2: What is the relationship between structural social inequality and personal wrongs?

Q3: What is the relationship between status harms and status wrongs?

I will put forward and defend a *status wrong account*. Identity-based wrongs and harms are distinctively and personally wrong because they involve a *status wrong*: in these cases, someone is treated in a way they shouldn't be treated. More specifically, identity-based mistreatment casts people as inferior and thereby fails to treat them as an equal.

My account is animated by the same ideas we find in the work of Chambers, Moreau, and in particular, Hellman. I argue that while these status-based accounts of the wrong of discrimination and identity-based harms are right to focus on status, they fail to capture the personal wrong incurred by people in these circumstances for two reasons. First, I argue that these accounts mistakenly rely on the power to inflict harms (chapter 2). We need to talk about status wrongs, not status harms. Second, I argue that Hellman's account rightly relies on interpretative value judgements, but that it lacks a convincing account of how these can be objectively valid and therefore can be taken seriously as a source of moral knowledge (chapter 3).

I trace both of these shortcomings back to an insufficient metaethics that is adopted by all of these authors, on which objectivity necessarily involves absence of all reference to subjectivity (chapter 3). Like most people working on these questions, these authors haven't thought much about the metaethical foundations of their views. And where they have, like Hellman, these foundations look shaky. If we want to maintain and develop the view that identity-based mistreatment is a distinctive wrong we need a metaethics that can underwrite it.

I spend the second half the thesis developing a Wittgensteinian metaethics that can underwrite the *status wrong account* of discrimination and identity-based harms and wrongs. I will do so by drawing on the work of Alice Crary. On this metaethics we can take interpretative value judgements seriously as a source of moral knowledge, without the need to appeal to further facts. This shift in metaethics allows me to rephrase the question. Instead of asking ‘*Why* are identity-based harms and wrongs distinctively and personally wrong?’, I can now ask: Does this intuition tell us something about how the world really is, or is it just a misleading sentimental response? I argue that this intuition tells us what the world is really like by drawing on and developing Crary’s metaethics (chapter 4). More specifically, I argue that authoritative perspectives allow us to see more, which is something that racist literature fails to do. Finally, I discuss how this meta-ethics can inform our thought on the meaning of social practices, taking affirmative action as an example (chapter 5). I argue that the dignity of a student depends not on whether they were admitted on an affirmative action programme, but on how they are treated by their university.

Chapter 1 – Are Identity–Based Harms Distinctively Wrong?

Some social groups are structurally privileged and others are structurally disadvantaged. Women, for example, are more likely to be a victim of sexual harassment, sexual assault and domestic violence, and to suffer from body insecurity, shame and a distorted sense of self–worth. Women are less able to make themselves heard or listened to across different contexts. And women are subject to a wage gap, an orgasm gap, and a glass ceiling. When women incur these or other harms associated with sexism, they often incur them because they are women. That is, they incur them in virtue of their social identity. I will call such harms *identity–based harms*.

The fact that people structurally incur these harms in virtue of their identity seems profoundly wrong. It is this intuition I explore in this chapter: Are identity–based harms distinctively wrong? If so, what feature of the world makes it so?

In this chapter, I take Clare Chambers' work on choice and gender–based injustice as a starting point (Chambers, 2004, 2008). Chambers presents us with two ways of vindicating the intuition that identity–based harms are distinctively wrong for the people incurring them. First, she argues that harms that depend on social norms are unjust. Call this claim the Social Condition. If it holds up, identity–based harms are distinctively wrong for the people incurring them because they are brought about by social norms. Second, she argues that harms distributed alongside social inequality cast people as inferior, and thus involve an additional status harm. Call this claim the Inequality Condition. If it holds up, identity–based harms are distinctively wrong for the people incurring them because they cast them as inferior and in doing so constitute an additional harm to someone's status. Both of these I take to be initially plausible. But, as I will argue in this chapter and the next, neither succeed in vindicating the intuition that identity–based harms are distinctively wrong.

I start by drawing out the intuitions I mean to discuss (1.1). After introducing Chambers' project (1.2), I argue that the reasons she presents for the Social Condition are insufficient to establish it (1.3). According to Chambers, harms that depend on social norms are unjust because they are contingent, and because it lies in our power to change these norms. I argue that contingency isn't sufficient to make a harm distinctively problematic, and that the argument overstates our ability to affect social change. I discuss an alternative reason to accept the Social Condition: we are responsible for harms that result from social norms (1.4). I argue that we're not straightforwardly morally responsible. Neither can other kinds of responsibility – such as political responsibility (Young) or a duty to end social inequality (Gardner) – help to vindicate the Social Condition, because the duty or responsibility to end structural problems doesn't translate to the level of a

personal wrong. I then move on to the Inequality Condition and the notion of status harm (1.5). According to Chambers, following norms that portray you as inferior violates equality, thereby constituting status harm. I will take up this idea in chapter 2.

1.1 Lottery-world

What is special about social inequality, compared to inequality based on other grounds? To get this question clearly into view, imagine a world that is like ours, except that a lottery determines who is privileged and who is disadvantaged.² This lottery marks people as ‘winners’ or ‘losers’, and rigs society accordingly. Just like in our world, some people accrue far more harms and disadvantages than others. For example, the harms involved in sexism described above will still often be incurred by the same person, while others won’t incur any of these harms. But unlike our world, *who* gets *what* isn’t based on social identity. There aren’t systematically worse-off social groups, only worse-off individuals. Is there something distinctively wrong about our world, compared to lottery world? The lottery world, surely, gives us much to regret. But while this is so, it seems intuitively plausible that there is something bad about our world that’s missing in the lottery world. Intuitively, there is something distinctively wrong about harms that people incur in virtue of their social identity, as opposed to the same harms incurred by brute luck. Intuitively, to be harmed is bad, but to be harmed because of your social identity is bad *and* unjust, unfair, or somehow morally regrettable. It is this intuition that I want to explore.

Drawing this comparison between lottery world and our world might seem unhelpful because, viewed in one way, the lottery world is exactly like ours. In our world, the way people acquire social identities is random. Most women end up as women by way of sheer luck (or perhaps sheer misfortune), much like they’ve taken part in a gender lottery. Once we understand this, the thought goes, we can see that the wrongs in both worlds are exactly the same: the harms and benefits are unequally and perhaps unjustly distributed by some lottery-like mechanism. In introducing this lottery world, we’ve merely introduced a copy of our world in which the same things are wrong for the same reasons. Reflecting on the lottery world won’t be of special help in understanding why identity-based harms could be distinctively wrong.

Viewed in another way, though, lottery world is light-years removed from ours. While it’s imaginable that people in lottery world still have social identities that on the surface somewhat resemble the ones that we have (e.g. woman), it’s unclear that they are in fact anything like ours.

² See *The Lottery in Babylon* in *The Aleph and Other Stories* (Borges, 1949).

In our world, being a woman dovetails with being disadvantaged. There are at least strong empirical links between being a woman and being disadvantaged in these specific ways, and perhaps even conceptual ones (see Haslanger, 2012). With gender entirely disconnected from gender-based harms, it is hard to imagine what being a woman in lottery world would look like. Moreover, while the way in which we acquire a gender identity might resemble a lottery, the way in which women subsequently incur sexist harms can't be captured by this simple lottery comparison. In our world, women incur these harms because of complex social structures and their history. But in the lottery world, we find none of these complex social dynamics at work. In this alien world, incurring these harms is simply a matter of luck. So, though people *ex hypothesi* incur the same harms in both worlds, the worlds are very different indeed. In our world, people incur these harms because of identity-based social dynamics. The question is what makes this mechanism different from a lottery, and, presumably, distinctively wrongful.

To aid our understanding, we can take this alien lottery world and slowly tune it to be more like ours. Imagine, for example, that those in charge of the lottery are thinking about changing the game, perhaps in order to make it more engaging. Instead of directly drawing lots for fortunes, there are three proposals on the table. On the first, people are first required to draw a colour lot. After this drawing, another drawing then determines which fortunes are attached to which colours. In a particular year, it could turn out that those with red lots are disadvantaged just like women are in our society. Those who drew yellow are luckier: they get the privileges we associate with being male. On the second proposal, people are required to draw colour lots whose meaning stays fixed. Red always marks out particular disadvantages and harms, and yellow always marks out certain privileges and benefits. On a third proposal, these fixed-meaning colour lots are hereditary. Whereas these proposals already looks more like the world we live in, it is a live question whether this minimal description of the mechanism is sufficient to generate the same intuitive wrongs concerning identity-based harms that intuitively seem to be there in our world.

1.2 Harm and unequal social norms

Chambers is interested in a specific kind of identity-based harms: harms that people choose for themselves in response to unequal social norms. In *Sex, Justice and Culture* (2008), she is particularly interested in harmful choices that women make in response to female beauty norms. Some women, for example, choose to have breast implants, wear heels, or cut their daughter's genitalia. What unites these choices is that they are not only physically harmful, but also that they are made in response to unequal social norms. According to Chambers, this makes them unjust.

Before diving into Chambers' account of what makes these choices unjust, it will be useful to briefly sketch the debate she positions herself in. Chambers is committed to a liberal political theory in which choice, autonomy, agency and freedom play a central role. Free choice keys us in to something valuable about human life, and should be respected and protected from interference. In choosing freely, people have normative power: they can turn an unjust situation into a just one by choosing it for themselves. Choice, that is, is a 'normative transformer' (Chambers, 2004). At the same time, Chambers recognises reasons to question this normative power, based on insights from social constructionism. Social constructionist theory recognises that free choice doesn't exist in a vacuum, and that all choice is situated in and thus constrained by its social context. Because of this, there are always further questions to be asked every time people choose to disadvantage themselves: How come they had the choice to harm themselves in the first place? How come certain reasons and not others were seen as valid? And how did they develop the preferences they based their choice on in the first place? Given that all choice is socially constructed and never an isolated act of agency, we shouldn't think of all choice as a normative transformer from the outset. The mere fact that these harms are chosen doesn't make them just. Instead, we should always take into account how a choice is socially constructed.

Chambers takes this tension between political liberal theory and social constructionism to be a productive one: learning about the ways in which choices are socially constructed can help us understand which choices are more morally problematic than others.³ Drawing on the work by social constructionist and feminist authors, Chambers (2004, 2008) argues that choice is unjust when:⁴

- 1) it is significantly harmful
- 2) the benefits of this choice depend on the acceptance of a social norm (the Social Condition)
- 3) *especially when* this social norm is unequal (the Inequality Condition)

The underlying idea is this: all choice is socially constructed, but some choices reflect and perpetuate systematic social inequalities. These choices are both the products and the driving forces of oppression, as they intersect with oppressive power dynamics. Even if they are *ex hypothesi* made freely and autonomously, there is still something deeply wrong about them: they are situated

³ Note that these choices are *ex hypothesi* free and autonomous: the question is whether there can be something unjust about a choice even when it is made freely and autonomously.

⁴ Chambers further argues that some unjust choices permit state intervention; see Laimann (2015) for a critical discussion.

in unequal social structures. This social inequality makes harmful choices morally regrettable over and above the harm they cause.

Put differently: harm is a bad thing to happen anyone, and so can be morally problematic, unless there is some circumstance that mitigates this moral problem. In general choice functions as a normative transformer. If people choose to be in pain themselves there is no moral wrong being done to them. But in these cases it doesn't, because society has put people in a position where it is beneficial for them to harm themselves and cast themselves as inferior, and this misfortune falls disproportionately on some people rather than others. It is unjust that these people incur harm, even though they chose to themselves.

Take, for example, breast implants. Breast implants are significantly harmful. On top of the health risks involved in any surgery there is a risk of serious lifelong health problems (Chambers, 2004, p. 23). These harms are severe, which prompts Chambers to investigate why this choice is there in the first place. And it doesn't look great. Unlike other kinds of surgery, the benefits of getting breast implants are purely dependent on social norms regarding female desirability. Without the social meaning we attach to large or perky breasts and the link between a woman's appearance and her worth, there wouldn't be a reason for women to get breast implants. To make things worse, the appearance norms that form the backdrop of this choice are deeply unequal. Only women are rewarded for harming themselves in response to these social norms, and their worth is tied to their appearance and ability to sexually gratify men. Given that breast implants are harmful and intrinsically tied up with social inequality, there is something distinctively wrong with choosing to get them, even if this choice is free and autonomous. Some choices, then, are unjust because their benefit depends on a social norm that offends against equality.

Chambers is interested in describing when and why these structural background conditions are unjust, and how they impact the function of choice as a normative transformer. I want to shift this focus slightly. I am interested in exploring how structural injustices like the one Chambers describes translate to the individual level. How does structural injustice translate to wrongs individual people incur in their lives as a result of living in these unjust conditions? For this purpose, I will use the language of individual wrongs. I will call choices people make under these conditions distinctively wrong. By saying these choices are distinctively wrong, I don't mean to say that these women make the wrong choice, act wrongly, or wrong themselves. Instead, I mean to say that the people that make these choices are morally wronged in some special way over and above the harm they incur; that something morally regrettable or bad or unjust is happening to them besides being harmed. There is a difference between a medically necessary procedure and

getting breast implants: the harm is the same, the difference is that there is some special wrong or injustice involved in the second. Or it seems like there is. I want to investigate that idea.

In the sections 1.3 and 1.4, I discuss the Social Condition ; in section 1.5 I discuss the Inequality Condition.

1.3 The Social Condition

According to the Social Condition, harming yourself in response to a social norm is distinctively wrong. Why? Chambers invites us to ask whether there are good reasons to follow the practice that don't depend on a social norm: if the norm weren't there, would it make sense for people to choose to harm themselves in this way? Sometimes, the answer is yes. It would still make sense to undergo surgery to remove a tumour: there are clear health benefits to it that don't depend on any social norm. There would also still be a reason to do harmful things such as smoke: whereas some benefits of smoking are social (e.g. socialising), others are physical (e.g. the nicotine rush) (Chambers, 2008, p. 197). Sometimes, however, the answer is no. Outside of our patriarchal norms that dictate female appearance, choosing to get breast implants simply doesn't make sense. It wouldn't make sense for women to get breast implants if it didn't have the meaning and benefit these norms give it.

Chambers is appealing to an intuitive thought: having to harm yourself *just* for social benefit seems wrong. But why exactly is that? She gives us two related considerations for accepting the Social Condition (2004, pp. 27–28, 2008, pp. 196–197, 265–266). First, when the benefit depends on a social norm, the harm is only contingently related to the benefit. The connection between the harm (e.g. getting implants) and the benefit (e.g. being more attractive or more normal) isn't a logically or conceptually necessary one. Instead, it fully depends on the norms we happen to endorse, in this case those of female desirability. Secondly, and relatedly, the social nature of norms means it lies within the scope of social action to change the norm. We ourselves keep this norm in existence, so we ourselves are able to dissolve it, and remove the benefits to making this harmful choice.

So, according to Chambers, choosing to do something harmful in order to gain a benefit that depends on a social norm is distinctively wrong because first, the harm involved is only contingently related to the benefits, and second, these social norms are changeable. In the Social Condition, then, we have found a first reason for thinking that identity-based harms are distinctively wrong. Much like choices that are made in response to social norms, identity-based harms in general depend on social norms and structures. When women are harmed in virtue of

their identity, the cause of these harms is social. Because these social structures are contingent and fall within the scope of social action, the identity-based harms they result in are distinctively wrong. This is why our world is worse than the lottery world, in which none of the harms depend on social norms.

I will now argue that whatever good reasons there might be for adopting something like the Social Condition, the reasons Chambers presents are insufficient to this end. Let's start with contingency. According to Chambers, when the benefit gained by a harmful choice depends on a social norm, the harm is only contingently related to the benefit. But this isn't necessarily the case. Consider, for example, a hormonal nasal spray that people with autism can use to become more sociable. Many people with autism don't regard their lack of sociability as a problem or obstacle – they regard it, rather, as society's problem with them.⁵ At least for some people with autism, then, the benefit of hormone therapy is purely dependent on social norms of interaction. But this harmful choice (interfering with your hormone levels) is in fact closely related to the benefit gained from it (alleviating deficits in social behaviour). So, benefits that depend on a social norm can also be intimately tied up with the chosen harm.

Moreover, there being a contingent connection between a benefit and a harmful choice isn't sufficient to make that harmful choice a distinctively wrong one. There are many choices that satisfy this condition which intuitively don't strike us as particularly problematic. Smoking, to take Chambers' example, is harmful because of the tar in cigarettes, but the presence of tar isn't necessarily related to the physical benefits of smoking, that are caused by nicotine. The same benefits can be had by taking nicotine some other way. Here, the contingent relation between the benefit and harm in smoking doesn't appear to be morally wrong. Contingency, then, isn't sufficient to make harms distinctively wrong.

We might think Chambers can grant that in some cases, the link between harm and benefit is contingent without it calling for moral outrage. There might rarely, if ever, be a necessary connection between a harm and a benefit. But this isn't a problem for Chambers, as what she is really worried about is cases in which the benefit's *very existence* is contingent, that is, dependent on social norms. This marks a difference between the harm in smoking from the harm in getting breast implants or using hormonal nasal spray: the benefits to the latter choices exist purely for social reasons, the benefits of the former choice don't.

⁵ <https://nos.nl/nieuwsuur/artikel/2283003-autistische-kinderen-mogelijk-socialer-door-neusspray-met-hormoon.html>

Contingency, then, marks a difference between these different kinds of benefits that come from harmful choices: those fully dependent on the social norms we happen to endorse and those that aren't. Granted this, why does contingency make these benefits distinctively wrong? One tempting thought is that in the cases of these thoroughly contingent benefits, the world very well 'could have been different', and that there is no 'real need' for people to make this harmful choice in order to gain the benefit. But, again, there are harms that fit this description that don't seem problematic. The mutation leading to cystic fibrosis 'could have never happened', and it doesn't 'need to exist'. Similarly, a perfect storm of factors leading to excessive damage and loss of life from natural disasters never 'need to exist'. The very existence of these kinds of harms are also contingent on some other factors being in place, but this kind of contingency doesn't seem wrong. Instead, the real worry seems to come in when the contingency is *social*.

What makes social contingency distinctively wrong? We can find the difference between mere contingency and social contingency if we reflect on how we use the word 'tragedy'. Consider the death of 39 Vietnamese migrants, who choked to death as they tried to make their way to the UK in a refrigerated truck in October 2019. Prime Minister Boris Johnson called this an "unimaginable tragedy".⁶ On the one hand, this seems to be an appropriate description of what has happened. On the other hand, however, it is jarring to hear a British politician like him describe these deaths as a tragedy. A tragedy is something that overcomes us, and that hits and overwhelms people much like a natural disaster. These migrant deaths, however, seem to be a direct and predictable consequence of British and European immigration policy. There is a tension between viewing something as a tragedy and having foreseen it while also having control over the circumstances that lead to it. These terrible deaths are *of our own making*: it's almost guaranteed these things will happen with our current policies.

So, we specifically mind these social contingencies because they are changeable *by us*. When we investigate Chambers' first reason for adopting the Social Condition, contingency, it turns out that it rests on the second reason, changeability: we are the authors of the social norms that these harms are based on, so it's in our power to change them.

Understood in one way, however, this second reason rests on a mischaracterisation of our power to change social norms. The mere fact that a cause of harm is social doesn't give us a reason to think it falls within our power to change it. Take, for example, the norms of female desirability and appearance that lead some women to get breast implants. These norms and the harms they result in exist in virtue of our thoughts and behaviour. But this social reality is real and complicated – a

⁶ <https://www.bbc.co.uk/news/uk-england-50150070>

lesson we keep learning from the feminist and social constructivist literature Chambers draws on – and it is persistent. Seeking to change the way in which individuals and groups behave and think involves disrupting deeply entrenched social structures, and such social change is notoriously hard to achieve. There is no reason to think that the social nature of this harm makes it easier for us to prevent it. If anything, the social nature of these harms presents a unique challenge to our ability to take away the cause of this harm. Changeability, then, can't account for thinking socially incurred harms are distinctively wrong.

1.4 Responsibility

1.4.1 Moral Responsibility

Perhaps we can defend Chambers' Social Condition in another way. Being changeable by us doesn't mark the morally relevant difference between identity-based harms and other harms. But even if social causes of harm are not easier to change, there is still an important fact about them that Chambers picks up on: we are the authors of these norms and policies. Because we author our social environment, the argument goes, we carry responsibility for it. If this social environment causes harm, that is on us. This responsibility makes identity-based harms distinctively wrong, because it makes us culpable (see Chambers, 2008, p. 211). The point isn't that we're in a great position to change problematic social norms, but that we *cause* them.

On this argument, socially caused harms are worse because we are morally responsible for them. What does it mean for someone to be morally responsible for something? We consider people morally responsible for something when they act. How exactly the responsible agent is related to her actions and the resulting harms is an ongoing debate. We might think that people are morally responsible for harms they could have prevented by acting differently, or harms that their action makes a significant contribution to. Alternatively, we might think that people are morally responsible for actions they did intentionally, or with awareness, or at least with negligence – they should have made sure to be aware of what they were doing. Or we might think people are morally responsible for the kinds of things that we deem blame- or praiseworthy. In all of these cases, the moral responsibility is held by individual agents.

The point here isn't to give an account of moral responsibility. Rather, the point is that on none of these accounts people seem morally responsible for the harms that result from social structures. People usually don't intentionally contribute to harmful social structures, and their individual

contributions are often negligible. Moreover, people's potential to affect social change and prevent harm by acting differently is small. We don't feel individual people are to blame or praise for authoring social structures, since how we act as individuals doesn't make a difference to these entrenched social norms. Social structures are upheld by groups of people, not by individuals. It might of course be sensible to say that those in power who design malicious policies or make grave mistakes that result in harm are morally responsible for this harm. But most socially incurred harms aren't like this. Most individual people most of the time don't seem to be culpable for harms that depend on social norms. So, even though a harm has a social cause, this isn't sufficient to make individual people morally responsible for that harm.

Perhaps my rejection of responsibility is too quick. While we don't straightforwardly cause social inequality as individuals, our behaviour is obviously in some way connected to these norms and the harms they result in. Surely there is more to say about how our individual behaviour is connected to structural injustice and how that bears upon our responsibility to act or refrain from acting in certain ways. These are two separate questions: First, if people's causal contribution to social inequality is negligible, what connection to a harm can make people responsible for it?⁷ Second, if our ability to affect change is minimal, why should we think harms that we are responsible for are distinctively wrong?

I will discuss two ways of answering these questions, arguing that neither are helpful at this point in the dialectic. The first idea is Iris Marion Young's notion of political responsibility which emphasises individual people's connection to structural injustice (1.4.2). We're causally linked to these injustices in *some* way, and this bears on our responsibility. The second idea is that we have a duty to change the world, from which flows a responsibility to act in some ways but refrain from acting in others (1.4.3). We can be culpable for acting, but also for failing to act where we have a duty to do so. While some aspects of this view are also found in Young, I will use John Gardner's view of discrimination to illustrate this idea.

1.4.2 Political Responsibility

Let's start with political responsibility (McKeown, 2018; Young, 2003). In putting forward this complementary notion of responsibility, Young is interested in explaining how it can be the case that people are responsible for participating in structural injustices, even though they aren't

⁷ This is a version of the collective impact problem: how can we say we're causally connected to an outcome if we don't make any difference to the outcome?

individually culpable for any of the wrongs that are brought about by the structures in question. Roughly, Young's idea is that when we're connected by our own actions to the processes that cause injustice for others, we have a forward-looking and outcome-oriented political responsibility for this injustice. The notion of political responsibility might be able to ground the Social Condition, in the following way. First, we are politically responsible for harm if this harm depends on social norms. Second, this political responsibility makes these harms distinctively wrong. So, harms that are the product of social norms are distinctively wrong.

First, then, what connection to a harm generates political responsibility? Young argues that people are politically responsible for structural processes in which they participate. Building on Young, McKeown (2018) argues that we are politically responsible for social structures when we are part of the process that causes them. People don't cause social structures to exist in the sense that their contribution makes a difference. Rather, our actions are *explanatory causal factors*: they are the background conditions that constitute structural injustice. More specifically, McKeown argues that we cause structural injustice by reproducing the background conditions in which we act. We do this by drawing on these background conditions when we try to bring about some state of affairs.

Take, for example, the way in which we cause the social structures that lead women to get breast implants. All kinds of people are trying to bring about some state of affairs: some women get breast implants to feel confident, some advertisers use women with large breasts to sell their products, and some men brag about sleeping with women with large breasts to boost their social status. In doing so, these people are reproducing the structural properties on which they draw for their actions. Their acts both draw on and reproduce the norms on which large breasts are sexy or beautiful, as well as the norms according to which a woman's value depends on her appearance. In reproducing these norms, the actors become politically responsible for them.

We should now be in a position to answer the second question. If individual people are politically responsible for the norms that they draw on and reproduce, why should we think the harms they result in are distinctively wrong? This, I believe, is where the notion of political responsibility fails to be of use. Consider the motivation behind the project of defining political responsibility. Intuitively, being connected to harms that result from social structures implicates us. We have the feeling that socially incurred harms are somehow worse. Political responsibility can be understood as an attempt to capture this intuition. It does this by describing in detail what connection to a harm or injustice we intuitively feel implicates us. But once we've defined political responsibility, it seems natural to ask why it matters whether a person is politically responsible for some injustice. Why does political responsibility show that the harms resulting from injustice involve a distinctive

wrong? This is where it seems like there's a subtle circularity going on. It's hard to see what answer we could give to that question that doesn't essentially involve just restating the initial ethical intuition that we're trying to vindicate: that harms that depend on social norms are somehow distinctively wrong. While giving an account of political responsibility might help us understand our causal connection to socially caused harms better, it doesn't help us understand why they would be distinctively wrong.

1.4.3 Duties or forward-looking responsibility

So far, I have talked about moral responsibility for our actions, and political responsibility in cases where our actions are causally related to an injustice. But there is another way to go: we can be culpable for acting, but also for failing to act. On this line of thought, we are sometimes responsible for failing to act to end unjust social structures. If we're responsible for failing to do something, there must have been a duty to do that thing in the first place. So, what we need is a positive account of our duties in the face of inequality and social injustice.

One such account of individual duties in the face of structural injustice is given by John Gardner (Gardner, 2018). Gardner starts by painting a picture of discrimination as a structural problem: the problem is that people with certain properties face discrimination in many aspects of their lives. Such structural disadvantage is unjust and we should end it. But, again, our individual actions alone won't change the world. And if our individual actions don't seem causally effective it's hard to see a reason not to discriminate, especially since non-discrimination can be costly. This, says Gardner, is a collective action problem: our society would be better off if no-one discriminated, but that only works if most people participate, and individuals have reasons not to participate. Gardner suggests that the state solves this collective action problem by assigning every individual the duty not to discriminate. This individual duty is derivative of the group duty to end the wider social problem of discrimination. Wrongful discrimination is wrong because it violates this duty.

But such an account by design cannot help us explain why harms induced by social inequality are distinctively wrong for the people incurring them. First, Gardner doesn't think of discrimination as inherently wrong. Instead, it is a *mala prohibita* – wrongful only when legally or socially made to be so (Gardner, 2018, p. 75). Second, and more importantly, Gardner denies that discrimination constitutes a personal wrong.⁸ On his account, discrimination is bad only because of its structural

⁸ Gardner leaves open the possibility that some cases of 'racist race discrimination' or 'sexist sex discrimination' constitute a personal wrong (Gardner, 2018, p. 78). It seems to me that all race discrimination is racist and all sex

effects. But such a structural account of the wrong of discrimination is deflationary: it denies that there is a personal wrong involved in discrimination and identity-based harms, and that those incurring these harms have a special complaint. Individual acts of discrimination aren't morally problematic; they are wrong only because of our legal duty not to perpetuate systematic inequalities.⁹

Other structural views of the wrong of discrimination have the same upshot. Glenn Loury's account of discrimination, for example, focuses explicitly on the social status of groups (Loury, 2002). Similarly, Owen Fiss suggests that discrimination is wrong insofar as it establishes or strengthens social castes (Fiss, 1976). On these accounts, the wrong of discrimination is analysed in terms of what it does to social groups and structures bigger than individuals. And so by design, they can't help us explain the personal wrong of discrimination, or more broadly, the personal wrong involved in identity-based harms.

It seems to me that on these structural accounts a part of the story is missing: how do these structural wrongdoings translate to the level of individual lives and interactions? We can fully agree that the wrong in identity-based harms depends on structural features of our world, and that we can't understand this wrong without understanding that it is structural. But we can still hope for more by way of an explanation of the wrongness involved in identity-based harms. That said, these structural accounts present us with a challenge: why think that there is a personal wrong of discrimination in the first place?

Let's take stock. Harms seem somehow worse when their existence depends on social structures: there seems something distinctively wrong about our world compared to lottery world. But these first attempts at vindicating this intuition by means of the Social Condition have run into serious problems, and Gardner warns us that this intuition might be mistaken. So, the question remains: how exactly should we understand the distinctive, personal wrong of discrimination? And why exactly should we think there is one?

1.5 The Inequality Condition and Status Harm

discrimination is sexist, meaning that they are tied up with racist and sexist social meanings, and I will go to argue that they all constitute a personal wrong.

⁹ Similarly, on Young's political responsibility, acts are wrong only because of our moral duty not to contribute to structural inequality.

Let's now turn to Chambers' Inequality Condition: there is something distinctively wrong about people harming themselves in response to a social norm when this norm is itself unequal. Maybe the problem is not so much that these harms are *social* – a claim that runs into trouble when we try to vindicate it. Instead, the problem is that these harms being based on social norms opens up room for them to be part of social inequality.

Chambers introduces the Inequality Condition as picking out a norm that requires some and not others to harm themselves. The problem, however, isn't just distributive inequality: the bare fact that some people or groups of people are affected and not others isn't enough to make a norm distinctively wrong. Instead, what is special about unequal social norms on Chambers' account is that they cast some people as inferior (Chambers, 2008, p. 210). How does this work? Chambers' main focus, as we saw, is gendered appearance norms that incentivise harmful choices. These norms sexually objectify women and predicate their worth on their ability to sexually gratify men with their bodies. They cast women as lacking power, weak, available, and submissive. In following these norms, women are portrayed as inferior to men. This offends against the equal moral status they have in virtue of being human.

Note that this appeal to equality thus depends on an account of what exactly it is about a specific norm that is degrading. Only those social norms that portray you as inferior will result in status harm if you follow them. This is something not all norms do: take, for example, the norms that regulate male appearance (2008, p. 210). Men who follow these norms are portrayed as independent, powerful, and strong. Following these norms gets them respect and status benefits. Even when men suffer serious other harms, such as the psychological and physical harms involved in complying with male appearance norms (e.g. taking steroids), they aren't degraded or cast as inferior in the same way that women are when they follow their respective appearance norms.¹⁰

Chambers equates being cast as inferior with a lowering of one's social status in a way that isn't compatible with moral equality. This offense to equality she calls a 'status harm' (2004, p. 13, 2008, p. 210). Such a status harm is to be understood as a kind of harm akin to, for example, physical and psychological harm. It can coincide with other harms (as is the case in breast implants) but can also exist independently (as might perhaps be the case with women wearing certain clothes or make-up).

This is a second reason for thinking that (some) identity-based harms are distinctively wrong. Identity-based harms are distinctively wrong when they involve an additional status harm, that is,

¹⁰ See *Force Majeure* by Ruben Östlund for an example of how adhering to masculine norms can harm men while simultaneously working toward retaining their superior status.

if the norm tied up with the harm casts you as inferior. The distinctive moral problem in our world seems to come from how these harms are distributed: they are incurred based on having a specific social identity that cast people as inferior. This way of distributing the harms creates status harm. Lotteries, on the other hand, are a uniquely ‘fair’ way of distributing harms or benefits. They are fair because apart from the unequal outcome in gains or losses, the distribution mechanism itself is indiscriminate and doesn’t impose any further status inequalities. Lotteries don’t stigmatize: no-one is worth less in virtue of losing a lottery, and no-one is worth more in virtue of winning it. Lotteries, then, don’t produce status harm. Moreover, in the absence of hierarchical social dynamics that determine who is worse off, it’s hard to even make sense of the idea of status harm. We don’t know much about lottery world, but we do know that it lacks many of the social structures that we find in our world. Without social norms to cast people as inferior, it becomes hard to see what it means to view them as inferior in the first place. Status harms are harms to our social status, and as such, are intrinsically tied up with structural social inequality. Without social inequality as the mechanism that brings them about and grounds them, we can’t fully capture the idea of status harms. In the lottery world, the thought is, status harms don’t exist.

So, we have an answer to our initial question: identity-based harms that work along the lines of social inequalities in our world produce additional status harm by casting someone as inferior, which is what makes our world distinctively wrong. But what exactly is the relationship between status harm and being cast as inferior? Can you cast someone as inferior without harming their status? And are symbolic affronts without status harm morally trivial? In chapter two, I approach the idea of status harm through legal and political philosophy on the wrong of discrimination.

1.6 Conclusion

Is our world of social inequality worse than a world in which the same harms are distributed unequally by a lottery (1.1)? Intuitively, the answer is yes. But what is distinctively wrong about identity-based harms? Based on Chambers’ account of gender based unjust choices, I explored two answers to this question. According to the Social Condition socially incurred harms are distinctively wrong for the person incurring them *because* they are social (1.2). I argued that the contingency and changeability of social norms aren’t good reasons to accept the Social Condition (1.3). Neither is an appeal to responsibility. First, it’s not clear that we’re morally responsible for harms induced by social norms (1.4.1). Second, political responsibility subtly assumes the validity of the intuition we want to vindicate (1.4.2). Third, accounts of duties in the face of social inequality merely understand the wrong on a structural level, not an individual one (1.4.3). According to the

Inequality Condition, harms are distinctively wrong for the person incurring them if they're incurred alongside social inequality (1.5). Chambers suggests that individuals are wronged by being cast inferior, and that we should conceive of this wrong as a harm to their social status.

Chapter 2 – Status Harms and Status Wrongs

From the work of Clare Chambers we got the suggestion that identity-based harms are distinctively wrong because they cast someone as inferior and as such amount to a *status harm*. The idea of status harm is developed in more detail by what I will call status-based theories of the wrong of discrimination. Status-based theories ground the wrong involved in discrimination in a violation of someone's equal moral status. On these views, discrimination is wrong because it subordinates, demeans, denigrates, or casts or marks people as inferior. Legal philosophers like Sophia Moreau and, I will argue, Deborah Hellman give an analytic reduction of the nature of these wrongs by conceiving them as (potential) harms in a way that is continuous with Chambers' notion of status harm. On the one hand, there is the social meaning of an act. On the other hand, there is what an act with that meaning *does* or *could do* to someone.

I believe these views are animated by the right idea: discrimination is wrong because it marks someone out as inferior. However, I argue that there are two serious problems with conceiving of identity-based wrongs as status harms. First, in response to Moreau, I argue that lowering someone's status isn't always wrong. Second, I argue that not all cases of wrongful discrimination actually involve status harm, and so cases of harmless discrimination pose false negatives.

Hellman's account of demeaning gives a more nuanced picture of the wrong involved in discrimination that aims to solve these problems. Addressing the first problem, she suggests that status harm is wrongful when it is tied up with the right social meaning. Addressing the second, she suggests that we focus not on the status harm per se, but on the power to inflict it. In this chapter, I will argue that the focus on status harm and the capacity to inflict it is a mistake as it mischaracterises the wrong of discrimination. If we are to take seriously the idea that the wrong of discrimination lies in casting people as inferior, we should let go of the idea that harms and potential harms are more real or more serious than expressive wrongs. In the next chapter, I will argue that this first move lacks a metaethical foundation, and attempt to provide one.

I start by discussing Moreau's account of unfair subordination and raising the first problem: lowering someone's status isn't always wrong (2.1). I raise the second problem in the form of cases of harmless discrimination that form false negatives (2.2). Hellman's account of demeaning aims to solve these problems by tying status harms closely to the right social meaning and focusing on the power to inflict status harms, rather than the status harm itself (2.3). Drawing on Elizabeth Anderson's work on relational equality, I discuss the relation between status harms and status wrongs and suggest that we should focus on status wrongs instead of (potential) status harms (2.4).

2.1 Moreau's Unfair Subordination

2.1.1 Subordination

On status-based theories discrimination is wrong because it fails to treat people as equals. 'Treating people as equals' is a widely used phrase that is equally widely attacked for being empty and meaningless (Moreau, 2020; Sangiovanni, 2017; Shin, 2009). What does 'equal treatment' mean? We don't think all people should be treated exactly alike. Nor is it helpful to say that all people who are alike should be treated alike: someone might simply claim that men and women aren't alike and that it's fine to treat them differently, and we'd have nothing to say to them based on our definition of equal treatment. The idea of equality seems too vague and too general to do the moral work we want it to: capture the moral difference between wrongful and permissible unequal treatment. It seems we need a story about what we *actually mean* when we appeal to equality. Moreau's account is an attempt to qualify what exactly is meant by status equality.

Moreau argues that discrimination is wrong when it *unfairly subordinates* (Moreau, 2019, 2020).¹¹ Roughly speaking, social subordination is the state of affairs in which one social group has a standing in society that is lower than that of another. Moreau specifies four morally salient features of social subordination:

- 1) Members of a subordinated group have less social and political power and de facto authority. They have a diminished capacity to do things with words, compel others, make themselves be listened to and be heard, and be taken seriously.
- 2) Members of such a group have or are ascribed traits that attract censure: negative moral attitudes are associated with you because of your group membership. In virtue of having some trait, you belong to a morally lower class of people, or are disposed to behave in a way that is worthless or even a vice. (e.g. 'Muslims are terrorists', 'women are manipulative')
- 3) Stereotypes work to forge the link between your group membership or trait and bad behaviour. Such stereotypes help to rationalize 1 – 2.
- 4) Structural accommodations tacitly overlook the needs and interests of the subordinated, which also helps to rationalize their subordinated position. These can be policies or practices, but also physical structures.

¹¹ Moreau describes two other ways in which discrimination can fail to treat people as moral equals: denial of deliberative freedom and barring access to basic goods.

Subordinating someone involves lowering their status or perpetuating their already inferior social status, characterised by 1 – 4. Subordination happens to individuals in virtue of their group membership. Discrimination can subordinate someone in two ways: either by playing a causal role in sustaining 1 – 4, or by directly constituting one of these aspects of subordination. Discrimination can, for example, constitute an expression of censure (2), or a structural accommodation that overlooks the subordinate's needs and interests (4).

Finally, marking someone as inferior is not just sending a message, but 'involves doing things in the world to that person'; 'altering their situation in certain ways' (Moreau, 2020, p. 72). Casting someone as inferior is problematic *because* it lowers your social status or perpetuates that lower status, i.e., subordinates you. It is this tangible setback of your social status that makes discrimination wrong.

2.1.2 What makes subordination unfair?

Merely lowering someone's status or perpetuating their lower status, however, isn't a pro tanto wrong. As Moreau concedes, subordination can sometimes be 'fair' or permissible (Moreau, 2020, p. 40). To illustrate this claim, she puts forward two pictures of justified subordination: a commanding officer and her subordinate soldiers, and a company manager and her subordinate employees. The soldiers and the employees stand in subordinate relationships to their superiors, the thought is, but they are not being subordinated wrongfully or unfairly. Such cases of subordination don't seem morally problematic in the same way discrimination based on protected traits – such as gender or race – is.

So when is subordination unfair? According to Moreau, the differences between these cases and wrongful subordination is that they are localised and based on the demands of a particular institutional role (Moreau, 2020, p. 41). In the case of the soldier and the employee, the differences in status are confined to a particular organisation, rather than across a whole number of social contexts. Moreover, this only seems fair insofar as these differences in power are justified by the demands of the job or role that the individuals have. In these cases, the differences in status can be justified by the needs of the army and a company, respectively. Moreau's position, then, is this: discrimination is wrong when and because it unfairly subordinates. Subordination is unfair, and thus morally problematic, when it happens to you across a wide range of social contexts. Discrimination based on gender, for example, is morally problematic because it subordinates women throughout their lives.

Now, does this account of discrimination as wrongful subordination help us explain how individual people are wronged when discriminated against? I'm not so sure.

First, the assumption that subordination is always problematic when it happens to you across a wide range of social contexts is false. Even widespread subordination can sometimes be justified and fair. Take, for example, the case of someone who is justly incarcerated. She has less social and political power, is stereotyped and censured, and her interests are often overlooked. She is, in other words, subordinated. This lower status permeates every aspect of her life. But her incarceration is 'fair' – it's required by her institutional role, and it's justified for the goal of the facility and justice system. Moreau wants to draw an important distinction between fair and unfair subordination. But as this example brings to light, she doesn't seem to have an independently plausible definition of fairness that allows her to distinguish between the two.

Second, Moreau's explanation is somewhat vacuous and circular. According to Moreau, discrimination is wrongful when it unfairly subordinates. Since subordination isn't always unfair, she has to specify what makes some subordination unfair. Her suggestion is that subordination based on a trait is unfair when it happens across a wide range of social contexts. But now we're right back where we started: it is wrong to discriminate against someone in circumstances of structural social inequality *because* it happens in circumstances of structural social inequality. Saying that subordination across a wide range of social contexts is unfair is, in a way, to bang your fist on the table and say: 'look, this issue is structural!' We do not get an explanation as to why that makes it problematic.

Third, it's unclear whether this approach is dialectically efficacious. I'm looking for a reason to think that discrimination is a personal wrong – that it wrongs individuals in a way that isn't completely derived from these structural effects. In my position, it doesn't seem dialectically efficacious to say: 'people experience discrimination many times in their lives, across all manner of different social contexts!'. If discrimination isn't a personal wrong to begin with, it doesn't help to point out that people experience this not-in-principle-problematic experience in many different places. Experiencing something that isn't a personal wrong many times doesn't make it a personal wrong. So, in my explanation of the personal wrong of discrimination it won't be helpful to merely refer to the fact that people get discriminated against over and over again. We want to know more about how this structural problem translates to the individual level.

What's missing is a story about what makes subordination wrongful. We started with the idea that wrongful discrimination casts people as inferior. Moreau tried to explain what is wrong about being cast inferior fully in terms of a setback to status. That is, she gave an analytic reduction of what it

means to violate equality. But that didn't work. We cannot capture the wrong of discrimination in terms of status harms or structural social inequality. It seems that we need equality to do *some* moral work.

2.2 Status Harm and Social Meaning

We need a better story to explain which instances of status harm or subordination are wrongful, and why. One *prima facie* plausible way of doing this is to tighten the link between status harm and social meaning. Lowering someone's social status is morally problematic when in doing so one offends against their equal standing as human beings. On this view, the wrong of discrimination is (a) a status harm combined with (b) the right expressive content – casting someone as inferior.¹²

This account makes progress on the problems I raised for Moreau in 2.1. Specifically, it explains why some instances of status harm or subordination are morally wrong: because they offend against equality. There is, however, a big problem with this view. Casting someone as inferior doesn't always come hand in hand with status harm. Sometimes an attempt at lowering someone's status is unsuccessful, or even backfires into accidentally raising someone's social status. Here is an example of such harmless discrimination:

Harmless Discrimination: Ben calls Betty a 'fucking bitch' in front of their colleagues. However, there is no up-take. Instead, all colleagues immediately defend Betty and tell Ben that he's being a misogynist.

In 'Harmless Discrimination', Ben doesn't inflict status harm. However, Ben seems to have wronged Betty just like he would have done if he had actually succeeded in lowering her social status. In discriminating against her, he seems to have wronged her independently of how things actually play out.

In cases of harmless discrimination no-one's status is actually lowered, and so according to the account under consideration there is no wrongdoing. But even when there is no harm, discrimination still seems just as wrong. Say that after the incident in 'harmless discrimination', Betty tells her friends about the way in which she's been maltreated. Simultaneously, Betty's boss disciplines Ben for his misconduct. We do not want to tell Betty and Betty's boss that they're being mistaken: that really, since Ben didn't manage to inflict any harm, they're wrong to think that he

¹² This view can be read into Chambers' *Sex, Culture and Justice* (Chambers, 2008). According to Chambers, casting someone as inferior constitutes a status harm.

did something morally objectionable. Moreover, it seems mistaken to let the wrong of discrimination depend on something as random as whether the discriminatory act is ‘successful’ or not. Even if he didn’t actually inflict status harm, what Ben did was intuitively still wrong.

We want an account that can capture the wrong of harmless discrimination. In the next section, I introduce Hellman’s account of demeaning. At first sight, Hellman seems to be able to capture harmless discrimination cases. However, I will argue that she ultimately isn’t successful in characterising the wrong we’re trying to capture.

2.3 Hellman’s Demeaning

2.3.1 *Demeaning*

Hellman puts forward another status-based theory of the wrong of discrimination on which discrimination is wrong when it *demeans*, and thereby fails to treat someone as a moral equal (Hellman, 2011, 2018). Demeaning has two aspects. First, there is the expressive dimension. A demeaning act or policy expresses that some person or some group is of lower status. Second, there is the power dimension. A demeaning actor or institution must have sufficient social power for this expression to lower someone’s social status.

First, let’s dive into the expressive dimension. Hellman characterises her view as a meaning-based view of wrongful discrimination (Hellman, 2018). Discrimination is wrong because it expresses that those discriminated against are inferior. People matter equally, and those expressions that deny this are morally problematic. Whether an act or policy expresses inferiority is a social fact. Racially segregated schools, buses and water fountains, for example, express that Black people are inferior. But segregating people based on something such as their last name does not. These cases are different because the history and current social status of a group affects what it means to treat someone who belongs to that group differently based on their social identity. Our cultural context and history invest what we do with meaning, and as such matters for the meaning of our actions.

But not all actions or policies that denigrate also demean. To demean, the person or institution that acts must also have the power to lower the social standing of the person they denigrate. Why is that? For Hellman, the answer lies in speech act theory as developed by Austin and Langton (Hellman, 2018, p. 103). As they point out, speech isn’t just about uttering meaningful words, but also about doing things with words: as we speak, we make requests, promises and apologies, and

issue orders. Such a speech act is only successful when a speaker is able to do things with words. This depends on the speaker's social status (Hellman, 2018, p. 103). I personally cannot, for example, order the Prime Minister to do something because I lack the relevant authority to do so. But ordering doesn't depend on actual effect: if the prime minister orders me to do something and I refuse, he has still given me an order. This is because he has the power to order me around. Demeaning, for Hellman, is such a speech act. In order to demean you need to (a) express that someone is inferior and (b) be able to lower their social standing.

Hellman makes progress on cases of harmless discrimination because she doesn't require an actual status harm for discrimination to be wrongful. Someone must have the power to lower someone's status, but it doesn't matter whether or not they actually succeed in doing so (Hellman, 2018, pp. 104–105). Hellman illustrates this aspect of her view with the Montgomery bus boycott. Prior to the boycott, Rosa Parks refused to give up her seat on the segregated bus to a white man even though the bus driver ordered her to sit in the back (Hellman, 2018, p. 104). According to Hellman, Parks was demeaned by the bus driver, even though in reality there was no social uptake and the demand actually resulted in lifting her social status to one of the US' most prominent civil rights activists. Similarly, Hellman can explain how Betty was wronged by Ben: she was demeaned by him because he had the power to lower her social status, even though as a matter of fact her status didn't get damaged.

Hellman's account of demeaning explains what makes a status harm wrongful by linking it to a denigrating social meaning. And it not only explains why actual cases of status harm are wrongful, but also gives an explanation of why near misses and unsuccessful attempts to inflict status harm are wrongful. As such, Hellman can make sense of cases of harmless discrimination like Betty's. What lies at the core of how these women are wronged is that these men have the capacity or power to lower their social status – whether they actually succeed in doing so or not.

2.3.2 The power to inflict harm

In this section I want to raise some difficulties for Hellman's view. First, let's get a bit clearer on how Hellman's view works in detail. According to Hellman, an act only demeans if the actor has the power or capacity to inflict status harm. What does it mean to say that someone has this power? According to Hellman, Ben has the power to demean Betty in virtue of his own social position – a male colleague, in this case. However, having such a social position doesn't automatically enable you to lower someone's status: in this particular situation, Ben wasn't able to do so. Moreover, a

high social position isn't necessary to do so either: you do not need a higher social position than your victim in order to damage their status. Slut-shaming, for example, can be done to women by women, and can be done as a way of *gaining* social power, rather than done from a position of power. And as we saw in chapter one, women can also inflict status harms on themselves. Social status seems neither necessary nor sufficient for being able to lower someone's status. If the power to lower someone's status derives from social status, we'll need some more detail on how social status gives people power to 'do things with words'.

Perhaps a counterfactual could help to explain how we can credit Ben with power: *had* Betty's colleagues not intervened, Ben would have lowered Betty's status. In virtue of that counterfactual, we can say that Ben did have the power to put Betty down even if he wasn't successful in this particular instance. This looks better. But we're not there yet. Compare Ben's case to another case in which I am (trying to) order the Prime Minister to do something. Do I have the power to do so? We can easily imagine a counterfactual world in which I am able to order the Prime Minister around. But that doesn't yet give me the power to do so in this world. So, it seems we need to qualify what counterfactuals count in giving an account of the power to do things with words.

What criteria could we set for the counterfactual case? Perhaps the counterfactual needs to be such that there is uptake. In Ben's case, *had* there been uptake, he would have successfully damaged Betty's status. But this criterion will always yield a trivially positive answer to any case: had there been uptake on my ordering the Prime Minister around, I would have successfully been able to do so, but that still doesn't make me powerful in this world. Saying 'he has power, because he would have successfully damaged her status had there been uptake' is just to imagine a situation in which someone did successfully do something with words, however far-fetched that situation might be. And this doesn't help us decide whether someone has the power to do things with words in our world.

What we want to say, I think, is that Ben has the power to put Betty down in virtue of living in a sexist society. We can explain why Ben has power by referring to these structural social inequalities. The relevant counterfactual situation has to be such that the social norms facilitating any potential status harm are unchanged. In the Ben and Betty case, we can imagine a situation in which Ben damages Betty's status in which the norms are unchanged. But in case of me ordering around the Prime Minister, the norms would have to be significantly different. So, we've ended up with the following view: in circumstances of structural social inequality, Ben has the power to put Betty down (even if he actually does not succeed in doing so, as in the case with the conscientious colleagues). If there are social norms which cast Betty as inferior, Ben has the power to lower her

social status. This also explains why women can slut-shame to gain social power, and why women can inflict status harm on themselves. The power flows from the social norms, not from their own social position.

I think this gets something important right: words are powerful because of this context of social inequality. But I also think Hellman's account gets something important wrong. Hellman intends to move away from a status harms view and tells us explicitly she is concerned with status *wrongs*, rather than status *harms*. As we saw, these wrongs don't depend on any harm actually occurring. Instead, she is interested in the speech acts, in what one *does* in discriminating against someone or treating them in a certain way. But this doesn't mean that there is no reference to harm at all. It just comes in one step further down, in Hellman's power dimension. According to Hellman, the demeaner needs have the power to damage someone's social status. On this account, the capacity to inflict status harm is still central in explaining why discrimination is wrong. In other words, the wrong of discrimination is still located in a potential harm or threat to someone's status. As Hellman insists on the power dimension of the wrong of discrimination, harm is still a fundamental and crucial part of her account. I will now argue that locating the wrong in a potential harm to someone's social status mischaracterises the wrong of discrimination.

To start with, the power bestowed on Ben by our sexist society can be limited by feminist work but that doesn't seem to mitigate the wrong. In 'Harmless Discrimination', it's not a fluke accident that Ben's attempt misfires – it's ensured that it does by the atmosphere Betty's colleagues have worked hard to put in place. Because of their feminist work, Ben has (much) less power than he might've had. But – and this is the point – this doesn't seem to matter. Betty doesn't seem any less wronged by Ben's comments simply because he might have had less power to inflict status damage. The way in which she has been wronged doesn't seem to depend on whether Ben's power to lower her social status is limited or not.

A defender of Hellman's view might insist that Ben has the unactualized power to lower Betty's social status, which does the heavy lifting in explaining how he wronged her. But this focus on unactualized power isn't very satisfying. In invoking it, we are referring again to the unjust background social conditions. But the existence and wrongness of those conditions isn't what's in question. What we're trying to capture is why Betty is wronged by what Ben does. And to say (in effect) that she's wronged because she's living in the kind of society where she is subordinated isn't to say anything about the wrongness of Ben's act in itself. Like in chapter one, we can agree that these unjust background social conditions are important, but still hope for more in an explanation of the personal wrong of discrimination.

It is tempting to agree with Hellman that the Betty case is a ‘near miss’, and that Ben tried to cause harm but failed. Or, side-lining Ben’s intentions for a moment, that Ben risked harm to Betty’s status but failed. This gives us an explanation of why Ben’s behaviour is condemnable: he has exposed her to risk. But this explanation mischaracterises Betty’s complaint. We want Betty to be able to say ‘he mistreated me’; ‘he treated me as if I was nothing’. On Hellman’s account we might cash_out this mistreatment in terms of threat or risk to one’s status: ‘he could’ve damaged my status’. But in doing so she locates the wrong somewhere else than where we first found it. The complaint isn’t that Ben risked her status damage – it isn’t about something that could have happened to her, or some risk that he exposed her to. It is about what actually happened to her in the first place. Ben has cast her as inferior, and this mistreatment by itself is a reason to complain. Betty probably wouldn’t, and shouldn’t need to, invoke some further hypothetical harm or risk that she has been exposed to in order to capture this serious mistreatment. Even in this ‘near miss’ case, Ben has actually wronged Betty in a way over and above risking status harm.

To drive home this point, consider a crude analogy to ‘Harmless Discrimination’:

Harmless Arson: Ben sets fire to Boris’ office, but Boris’ colleagues douse the fire before it causes any harm.

On Hellman’s story, Ben has wronged Betty in much the same way as he is wronging Boris: by risking them a serious injury. But this case seems importantly different from discrimination cases. It seems that in case of the fire, the worst is actually prevented by preventing the harm. If there are repercussions for Ben, it will be because he tried to harm Boris or risked harming him. But in cases of discrimination, it seems that the worst thing about the whole situation is the fact that Ben called her a ‘fucking bitch’ – regardless of subsequent harmful effects this might also have, or could have had. Any repercussions Ben faces would be legitimate because of what he has actually done, not because of the harm he has risked. Harmless arson is a failed attempt at arson. But harmless discrimination isn’t like this. Harmless discrimination isn’t a failed attempt at discrimination, it is just discrimination.

Capturing the wrong done in terms of risking harm to one’s status seems to mischaracterise the wrong and locate it in the wrong place. I want to suggest a better way to understand the wrong of discrimination is as follows: Ben is treating Betty in a way she shouldn’t be treated. Betty is being wronged because she has an interest not being thought of or spoken of as lower than others. This interest doesn’t derive from the way in which her social status is affected by Ben’s act. Ben is wronging her in a way that isn’t appropriately cashed out in terms of a threat to her social status.

2.4 Harms and Wrongs

2.4.1 *Tangible Harms*

I've suggested that if we're going to move beyond a deflationary view of the wrong of discrimination that we get from Gardner, Loury and Fiss, we should be talking about status wrongs, not status harms. But this is exactly the kind of language Chambers, Moreau, and Hellman use to motivate their accounts. So why do they feel the need to cash out the wrong in terms of status harms?

For Hellman, it is important that demeaning is something that you do in the public world, like ordering. When you demean someone, you are not merely expressing they are of lower worth, you are doing something to them that can be publicly appraised and could make a tangible difference to their lives. And so even in 'near miss' cases and cases of harmless discrimination, the wrong is intimately tied up with the power to make this tangible difference to someone's social status. In a similar vein, Moreau and Chambers emphasise that discrimination and identity-based harms are distinctively wrong because they involve a tangible setback of one's social status.

If we're to give a story about the wrong of discrimination, Hellman and the others seem to say, we need to give a story on which casting someone as inferior is tied to the publicly appraisable facts about their social status. Being treated as lesser than others isn't some free-floating, harmless issue. It's not just about expressing that someone is inferior, or using degrading language. It's a serious wrong. To appreciate its seriousness, we should recognise that what discrimination does is real, and changes people's lives for the worse in a real, objective sense – where 'real' and 'objective' refers to publicly appraisable, tangible facts. To appreciate the seriousness of the wrong of discrimination, we need an analytic reduction of the wrong involved. We must link it to something tangible, like status harm. If we don't, we end up with a free-floating, vague appeal to equality.

This assumption underlies status harm accounts and structural views of discrimination alike. I will now suggest that this assumption gets things the wrong way around. We do not need to appeal to status harms in order to be able to take status wrongs seriously. Instead, the wrong of discrimination starts with the appeal to equality: people have an interest not to be treated as lower than others, and treating them as such amounts to a status wrong. Status harms mostly matter in light of status wrongs, not the other way around. In order to make this point, I will first discuss

the relation between status inequality and material inequality in 2.4.2, drawing on Dr Seuss and Elizabeth Anderson. I will return to the relation between status wrongs and status harms in 2.4.3.

2.4.2 *The Sneetches*

What is the relationship between status inequality and material inequality? To get this question into view, I want to turn to a world of fiction presented to us by Dr Seuss: the world of the Sneetches.¹³ The Sneetches are yellow bird-like creatures split into two social classes: the privileged star-bellied Sneetches and the inferior plain-bellied Sneetches. The star-bellied Sneetches brag that they're 'the best kind of Sneetches', refuse to play and eat with the plain-bellied Sneetches, and walk past them 'with their snoots in the air'. Then a businessman appears. Looking for a way to profit off the Sneetches' social class system he offers the plain-bellied Sneetches the use of his Star-On machine, at a cost of \$3. The treatment is instantly popular, with all the non-starred-Sneetches lining up to get stars. This greatly upsets the original star-bellied Sneetches who are now in danger of losing their special status. The businessman then wheels out his Star-Off machine and charges the original star-bellied Sneetches \$10 to take their stars off in order to be exclusive again. But the newly starred Sneetches buy this treatment too. This escalates, with all Sneetches frantically running through both machines – star on, star off, star on, star off – until finally all their money is spent. The businessman packs up his machines and leaves a rich man. From this debacle, the Sneetches have learned that neither group is superior to the other. Once they realise this, status inequality disappears. They forget all about stars, and all play and eat together.

In the hands of Dr Seuss, the Sneetches are painted as incredibly silly and status-obsessed creatures. If the Sneetches would simply stop caring about status so much, the message is, everything will be fine and status inequality will disappear. Status inequality isn't a *real* problem; there aren't *real* differences between the Sneetches, and once everyone recognises that the problem dissolves. Our world isn't like this. The story is funny and the Sneetches are silly only because their status inequality isn't really hooked up to any real or tangible inequality and harm. In our world, status inequality seems a lot less silly. Having a lower status can make your life go worse for you in very real material ways: in terms of work, housing, health, safety, or self-image, to name a few. For human beings, status inequality is closely tied to material inequality. As such it seems we need to draw on this tangible, material structural inequality and injustice to explain the wrong of status discrimination. At the same time, it seems it would be a mistake to say that all there is to status

¹³ <https://www.youtube.com/watch?v=PdLPe7XjdKc>

equality is material equality. It's not the case that status only matters insofar as it means you have less access to good things, or insofar as it means that bad things happen to you. Social status matters its own right, too: we don't want to tell people just to stop being silly and stop caring about social status so much. Being treated as equals is important to human beings.

This narrative suggests that the challenge for theories of the wrong of discrimination is this: to tie the wrong of discrimination to the material circumstances tightly enough to explain its seriousness (to avoid making discrimination look silly, as with the Sneetches), but not so tight that status itself isn't morally important anymore (to avoid a picture on which discrimination isn't morally problematic in itself, but only insofar as it contributes to material inequality). Material inequality needs to be part of the story to account for the seriousness of the wrong, but also cannot be the whole story. So, the challenge seems to be to explain how status inequality is tied to material inequality – not too tightly, but not too loosely either.

Elizabeth Anderson turns this way of thinking upside down in her paper *What Is The Point of Equality?* (Anderson, 1999). Anderson critiques egalitarian thinkers for focusing too much on distribution of material goods. Why, she asks, does this material distribution matter in the first place? Her answer is that it matters insofar as it shapes the conditions for relational or democratic equality: people have to be able to meet as equals, and be regarded as equals. A materially equal society could still be relationally unequal. In such a society, respect isn't equally distributed. Of course, you cannot just distribute equal respect as easily as you could distribute money. But you can make the material distribution so that everyone has enough material resources to be able to look other people in the eye and engage them as a social equal, and to escape relations of domination and unequal respect.

This suggests a different picture of how status inequality and material inequality are related. Status inequality is important in its own right, and we don't need to refer to material inequality in order to take it seriously. It is, if anything, the other way around: we need to refer to status inequality in order to explain why we should take material inequalities seriously. Material equality often matters *in light of* status equality. Dr Seuss was wrong to think that status inequality is silly. But the way to fix that picture isn't to insist that status harms are intimately related to material harms. Instead, we should insist that status equality is important in its own right. Being recognised as equals just is important to us as human beings.

2.4.3 Status Harm and Status Wrong

Material inequality often matters *in light of* status equality. But, as we have seen, we can understand status inequality in two distinct ways: status wrongs on the one hand, and status harms on the other. A status *wrong* involves treating someone in a way they shouldn't be treated – by casting them as inferior, treating them as lower, denigrating them, and the like. This is the expressive dimension of discrimination. A status *harm* involves lowering someone's social status or perpetuating their lower social status. This is the tangible dimension of discrimination. So what is the relation between status harms and status wrongs?

Status-based theories of the wrong of discrimination seem to face a dilemma similar to that described in 2.4.2. Discrimination is highly correlated with tangible status inequalities and subordination. Any account of the relevant wrong that totally brackets that off, and tries to locate the wrong entirely in some kind of expressive dimension, seems to be missing something. At the same time, it seems it would be a mistake to suggest that all there is to status wrong is a status harm. Any account of the relevant wrong that fundamentally grounds the wrong in the status harms that are ordinarily correlated with discrimination has a counterintuitively deflationary upshot – as we saw in 2.3, it would fail to capture the wrong in cases of harmless discrimination.

Hellman appeals to the power to inflict status harms in order to find the sweet spot between Dr Seuss' view (status is just silly because it's not tangible) and Moreau's and Chambers' view (status matters because it is tangible). She tries to find a way of tying the status wrongs closely enough to status harms to explain its seriousness, but not so tight that status wrongs aren't important in their own right anymore. In doing so, she accepts a premise of the conceptual space these views are in: in order for status wrongs to matter morally, we must tie it back to the tangible reality of status harms. But this, I think, gets things the wrong way around. Status harms matter morally *in light of* status wrongs. We've already seen two pieces of evidence for this. As I argued in 2.1, a mere status harm isn't wrongful. We need to appeal to the expressive dimension of discrimination in order to explain why some ways of occupying a lower social status are morally problematic. So, tangible status inequality matters insofar as they feed expressive status inequalities, not the other way around. And as I argued in 2.3, a status wrong view offers a more direct and faithful interpretation of intuitive moral judgements about all the cases the account is supposed to cover. Finally, there is no real motivation for accepting the status harm account unless you assume that it's trivial to worry about moral wrongs that are materially harmless. But this assumption is somewhat question-begging. It seems to be borne of a mistrust of intuitions, and there's no way to justify that mistrust without appealing to a premise that already expresses that very same mistrust.

Hellman, Moreau and Chambers are interested in status wrongs, but seem to feel disentitled to talk about status wrongs in their own right. I've suggested that this is mistaken. We can and do care about status wrongs – violations of equality – in their own right. We do not need to appeal to status harms in order to be able to take status wrongs seriously. Instead, understanding the wrong of discrimination starts with the appeal to equality: people have an interest not to be treated as lower than others, and treating them as such amounts to a status wrong. Status wrongs are serious in and of themselves.

2.5 Conclusion

Why is discrimination wrong? On status-based theories discrimination is wrong because it fails to treat people as equals. How should we understand this claim? I've argued that there are serious problems with conceiving of these status wrongs as (potential) status harms. First, lowering someone's status isn't always wrong (2.1). Second, not all cases of wrongful discrimination actually involve status harm (2.2). Third, the focus on status harm or a capacity to inflict status harm mischaracterises the wrong (2.3). I argued that the felt need to ground status wrongs in (potential) tangible harms comes from the mistaken assumption that we need status harms in order to take status wrongs seriously (2.4). Instead, I argued that we can take status wrongs seriously on their own accord.

Chapter 3 – Subjectivity and Objectivity

Why is it especially bad to be worse off in virtue of your social identity? In chapter two I tried to answer this question by appealing to status-based theories of the wrong of discrimination. I argued that we should conceive of the wrong of discrimination as a status wrong, rather than a (risk of) status harm. Human beings have an interest in being treated equally that is violated in discrimination. In this chapter I outline a metaethics that can underwrite the status wrong account.

The status wrong account faces two explanatory challenges. First, what does treating people equally actually amount to? In chapter two I argued that we need our intuitive grasp of equality to do real moral work, since we can't just give an analytic reduction of what failing to treat others as equal look like by means of status harms. (2.1). But this view is often met with the worry that the appeal to equal treatment is too general and too vague to be able to distinguish between wrongful discrimination and morally permissible treatment (2.1).

Second, why should we take status wrongs seriously? Proponents of structural theories of discrimination deny there is a personal wrong involved in discrimination (1.4.3). I've since argued that a status harm account can't successfully capture this wrong (chapter 2). Instead, we should look to status wrongs. But while we might in general grant that you can wrong someone without inflicting harm, it's not yet clear why we should think that these status wrongs are *real* if we can't conceive of them as potential harms. What philosophical or moral foundations do we have for thinking that status wrongs are really wrongs? Faced with this question, merely positing an interest in being treated equally seems ad hoc.

Building on Hellman's work, I will argue that we can have true judgements about whether people are being treated as inferior in a morally problematic way that offends against equality. But as it stands, Hellman's account is metaethically undercommitted. Roughly, the problem is this. Interpretative value judgements are subjective judgements. They are necessarily made from a particular point of view and depend on subjective ways of knowing. But in order for us to take them seriously, we want to say they are also objective: they capture what things are really like independently from any point of view. These judgements are judgements about how things *seem* to us, so how could we reasonably disagree about them?

Hellman accepts this picture of subjectivity and objectivity as mutually exclusive and suggests interpretative value judgements lie somewhere in between: they are 'modestly' objective. I argue that her account fails because even though it recognises the need for subjective ways of knowing, it still subscribes to a picture of objectivity as a 'view from nowhere'. Based on the work of Alice

Crary, I put forward an alternative metaethics that can underwrite interpretative value judgements. On this metaethics, objectivity doesn't require excluding all reference to our subjectivity. Instead, our non-neutral point of view is the only way to access certain ethical facts. If we want to have a meaningful discussion about what it is to be human and how human beings merit to be treated, we can't adopt a neutral stance as to whether or not discrimination is wrongful. It's implicit in the ethically inflected way that we perceive human beings that the judgements that we make about these cases are sound.

I begin by discussing the questions of vagueness and seriousness (3.1). In response to the question of seriousness, I outline Hellman's account of modest objectivity and argue that it is insufficient to ground interpretative value judgements (3.2). I then draw on the work of Crary to argue that interpretative value judgements can have objective status (3.4). Changing the metaethical approach will allow me to rephrase what kind of story would allow us to accept this intuition as morally informative. Rather than looking for a further fact to ground the intuition, we can now ask whether the intuition tells us something about how the world really is. In this chapter and the next I suggest that it tells us what the world is really like.

3.1 Interpretative value judgements

According to Hellman we can have true judgements about when people are being treated as inferior in a morally problematic way that offends against equality. These judgements can at once be interpretative (requiring our subjective engagement), value-laden (moral), and objectively true (they tell us something about what the world is really like). In this section, I will discuss the question of vagueness and the question of seriousness.

3.1.1 Vagueness

How does the appeal to equality help us to differentiate wrongful discrimination from other kinds of unequal treatment? On moral equality-based views such as Hellman's, discrimination is wrong when and because it fails to treat those affected as moral equals. These views assume we already have some intuitive grasp of what it is to violate someone's equal moral status. But, Sangiovanni argues, this intuitive grasp isn't specific enough (Sangiovanni, 2017, p. 131). It doesn't tell us in what sense and why human beings are of equal worth, and what constraints on our behaviour this is supposed to set.

Sangiovanni motivates his concern with the case of in-work benefits for EU migrants. In 2016, the UK government sought to restrict the access of EU migrants to in-work benefits for the first four years of residence. This is discriminatory in some sense – it distinguishes people based on their migration status. But is it morally problematic? Hellman's account, Sangiovanni argues, leaves us without tools to address this problem. Hellman reframes the question in terms of equal worth: from 'Is this policy wrongfully discriminatory?' we get 'Are these people treated as being of lesser worth?'. But this doesn't get us closer to an answer. Hellman doesn't tell us what moral equality entails, so we still don't have any criteria for determining whether these people are being treated as inferiors in a morally problematic sense. Hellman's appeal to equality, Sangiovanni contends, is seriously underdetermined (Sangiovanni, 2017, p. 131).

If this is Sangiovanni's criticism, it seems to miss its target. Hellman doesn't just suggest we reframe the question in terms of equality, but also gives instructions about how to go about answering that question: we need to make an interpretative judgement about the social meaning of an act (2011, p. 61). In doing so, we should use our interpretative and empathetic skills to pay attention to the speaker, context and the content of the words used.

Let's try to do just that. As a general background, in-work benefits are paid to people who are working a certain number of hours a week but don't meet an income threshold. These benefits exist because without them people fall below a certain standard of living into poverty. People on benefits are highly stereotyped in negative ways. They are stigmatized as 'welfare mums' or 'welfare queens', 'lazy', and substance abusers. On top of this, the working poor with EU citizenship also fall victim to xenophobic, racist and dehumanizing stereotypes. Romanian, Bulgarian and Polish working poor, for example, are viewed as 'scroungers' and foreigners who come to live comfortably in British jobs, British money and British healthcare that they are not entitled to. They are said to weaken the UK from the inside, and, as evident from the Brexit referendum, many want to 'take back control' by stopping them from working in the UK. In addition to this hostile environment, the jobs these EU working poor work are looked down upon, offer bad pay and bad working and living conditions. Finally, cutting in-work benefits for this group would have low monetary gain for the UK government.¹⁴

From the context, it becomes clear that the social meaning of this policy is xenophobic, racist and classist. It singles out people for worse treatment based not on any financial concerns, but based on their inferior social status in British society. This signals that the policy is one that draws on

¹⁴ <https://migrationobservatory.ox.ac.uk/resources/reports/eu-migration-welfare-benefits-and-eu-membership/>

and legitimises these xenophobic, classist stereotypes. Attending to the empirical details of the cases makes clear that the policy casts these people as inferior in a morally problematic way.

Hellman runs through various examples just like this in her book. Sangiovanni, then, seems to judge too quickly. If we take Hellman's recommendations seriously, her moral–equality based view of the wrong of discrimination does provide the tools to talk about being cast as inferior in a specific way. If we consider potential cases of morally problematic unequal treatment and attend closely, we're able to judge which ones are problematic and which ones aren't.

3.1.2 *Seriousness*

If there is a point to Sangiovanni's objection it can't be that the account lacks detail, nor that it lacks a procedure for determining whether a certain treatment or policy is morally problematic. Rather, his problem might lie with the *kind* of procedure that is offered. Sangiovanni is asking how we could deem a policy to be wrongfully discriminatory if an interpretative value judgement is all we have. There is, he might want to say, no compelling reason for him to agree with this analysis other than that it might *feel* right. But, he might continue, feelings and opinions alone can't make up the robust kind of moral judgement that we need when we want to condemn discrimination. The point can't just be that we learn from our emotional and moral sensitivities that this kind of thing is wrong. Because if that is so, anyone who were so inclined could just disagree, and we would have nothing to say to them. This would spiral into both parties digging their heels into the sand and saying 'I am right, because I *feel* I am'. In this right/wrong game, there seems to be no way to appeal to independent reality, no way of tracking whether our feelings are the right ones, and thus no possibility for reasonable disagreement. We're just describing how things *seem* to us, not disputing what things are really like, so all we're left with is a mere difference of opinion. We've now landed at the second question: what's to say we can take these judgements seriously?

The argumentative strategy pursued by the authors in chapter one and two is to find some further fact to ground this emotional moral judgement in. It seems we can't have substantive disagreement if the final court of appeal is how the facts seem to you, rather than the facts themselves. As such, the thought is, we should look for further facts that make discrimination wrong and that can back up our interpretative value judgement. But in chapter two I argued that this strategy doesn't work. First, value–neutral accounts of the wrong of discrimination aren't able to distinguish between morally problematic subordination and permissible subordination (2.1). So, I argued, we need this value judgement to be there from the start: particular kinds of subordination just are bad for human

beings. Second, grounding the wrong of discrimination in a tangible fact leads to false negatives (2.2). And third, doing so mischaracterises the personal wrong of discrimination (2.3).

So, what other reason could we have for taking these interpretative value judgements seriously? In response to this second challenge, what I want to say is this. This judgement by itself is sufficient to morally condemn discriminatory behaviour, acts or policies. Once we make this judgement based on the appropriate evidence, there is no need to refer to further facts that makes it so. These moral judgements can be objectively valid: they capture what things are really like. As such, we can take them seriously without appealing to further facts. If we ignore them and try to avoid leaning on subjective modes of understanding we deprive ourselves of an important source of moral information. This is the response I aim to vindicate in this chapter. This requires metaethical work. In 3.2, I discuss Hellman's attempt to achieve 'modest objectivity' for these judgements. I will reject it and offer an alternative metaethics to underwrite her account and support its wider argumentative purposes.

3.2 Modest Objectivity

3.2.1 *Modest Objectivity and Ideal Epistemic Conditions*

We want to say that we can take interpretative value judgements seriously because they are objectively valid and tell us what the world is really like. But this requires some metaethical work. These judgements are judgements about how things *seem* to us, so how could we reasonably disagree about them?

Hellman develops some terminology to capture this problem. On what she calls *strong objectivity*, what is the case about the world never depends on what humans take there to be (2011, p. 74):

A fact is strongly objective if it is true or false independently of anyone's reaction to it or perception of it. It would be true or false even if there were no people in the world at all. It is simply the way things are.

Strong objectivity, Hellman goes on, is 'inappropriate for interpretative judgements of social practices precisely because they are judgements about social practices and therefore involve us' (2011, p. 74). Demeaning depends on conventional practices and norms for showing disrespect, and thus depends on people's interpretation of the situation. So, the fact that an act demeans cannot be strongly objective.

But neither do we want to say that these judgements are what Hellman calls *minimally objective*. On *minimal objectivity*, what is the case about the world depends on a majority agreement (2011, p. 72). Minimal objectivity would leave us unable to challenge dominant but false views, and would instead relegate the issue to a mere matter of taste. Instead, we want to be able to have reasonable disagreement. We want to say that there is a moral fact of the matter as to whether a particular act or practice demeans, and that the majority can get it wrong.

Having set up this contrast between strong and minimal objectivity, Hellman suggests that we look for a sensible middle ground: interpretative value judgements can be *modestly objective*. On *modest objectivity*, what seems right under ‘ideal epistemic conditions’ by ‘ideal judges’ determines what is right (2011, p. 75). A familiar example of modest objectivity is colour judgements. Colours don’t exist independently of our perception of them, but they are ‘out there’ in the world, and judgements about them can be true or false. It’s possible for all of us to get them wrong, and there can be reasonable disagreement about them, instead of a mere difference of opinions. In this sense, they are objective. According to Hellman, facts about whether someone is demeaned are like facts about colour. These facts don’t exist independently of our perception of them: their existence depends on what we take the world to be like and we can only access them through subjective methods. But judgements about them can still be true or false.

Hellman doesn’t aim to put forward a conclusive account of what these ideal judges and ideal epistemic conditions look like. But she does give us some suggestions that I will briefly outline here. One line of thought Hellman is interested in exploring is maximising empathy. By cultivating empathy for another’s point of view and using our imaginative skills, the thought is, we put ourselves in a position from which we’re able to perceive the relevant facts. However, Hellman points out, maximising empathy alone isn’t going to get us where we need to be. Uncritically adopting one particular perspective (the potential victim’s) and accepting any claim they make as legitimate, whatever its merits, doesn’t get us an objective judgement. Doing so would imply that anyone who claims to be aggrieved is always right. And that is the wrong outcome.

Hellman’s rough sense of how to navigate this is that we’re to pay special attention to the perspective of the aggrieved and then make an ‘independent assessment’ of the value of their perspective. She frames this question in terms of impartiality: we need to use (only) the relevant and appropriate criteria to come to our value judgement. In order to do so, we need to specify relevant and irrelevant criteria for these judgements. Hellman makes a start on this by discussing one relevant and one irrelevant criterion. Amongst the relevant criteria is the subjective experience of those acted upon – hence the importance of the victim’s own sense of how she’s being wronged.

Amongst the irrelevant ones is personal like or dislike: the ideal judge must be free of bias. Recognising that this is only a rough starting point, she writes: ‘As we learn more about which practices demean (through conversations with others, empathetic thought experiments, etc.), we will come to have a clearer sense of what the ideal epistemic conditions are for making such judgements’ (2011, p. 78).

3.2.2 Objectivity and Subjectivity

I will now argue that this sensible middle ground isn’t so sensible, as Hellman is trying to reconcile an understanding of objectivity and subjectivity that are fundamentally opposed.

The problem with classifying interpretative value judgments as strongly objective, I take it, is not that facts about whether a policy is demeaning or not wouldn’t be true or false if there were no people in the world at all: facts about the average height or GDP wouldn’t be true or false if there were no people, but that doesn’t seem to bar them from being objective in a strong sense. The problem isn’t that these facts are about human beings and their social practices, or even that these facts depend on how people interpret things. Rather, the problem is that interpretative value judgements are *subjective*, and that we’re used to characterising objectivity in contrast to subjectivity.

In mainstream philosophy, objectivity and subjectivity are contrasted in at least two ways:

- 1) Subjective judgements are relative to a point of view, whereas objective ones are unrestricted by any such points of view. Objective judgements should – at least in principle – be ‘made from nowhere’, or at least we should be able to make them ‘from anywhere’.
- 2) Subjective judgements are sensational responses, whereas objective judgements represent what there is independently of how we take it to be. Subjective facts are facts about how things seem to us, objective facts are what things are really there.

In other words, we’re used to thinking about objective facts as accessible through neutral, non–subjective methods 1) not from a particular point of view and 2) being the case independent of what the one casting judgement takes to be the case.

Interpretative value judgements are fundamentally subjective in these two ways: there is no way to get such knowledge while not depending on our subjective ways of knowing things, and what is ultimately the right judgement depends on what human beings take to be the case. First, interpretation requires that we approach a culture ‘from the inside’: we have to think in the terms of that culture if we’re going to understand the social meaning of an act or policy and whether it

is morally problematic. This cannot be done from a ‘nowhere’ or ‘anywhere’: interpretation is done from within a culture, and thus depends on our subjective and necessarily parochial ways of understanding the world, and our take of how things are. Second, because these judgements are moral judgements, they depend on a value-laden way of looking at the world: we can only make such judgements if we already know what morally matters in the lives of human beings and what kinds of treatments are bad for them. While it seems we are able to reject particular takes of how things are, it’s not possible to avoid the human perspective altogether. Third, our ability to interpret the social meaning of an act heavily relies on our subjective capacities. Hellman, for example, emphasises the need for empathy; our ability to relate to others and put ourselves in their shoes, and imagine what things might be like for them. So, interpretative value judgements are subjective: they are accessible only through subjective methods 1) from a particular point of view and 2) depending on what the one casting the judgement takes to be the case using their subjective ways of engaging with the world.

Recognising this, Hellman wants to resolve the tension between subjectivity and objectivity by aiming for the sensible middle ground. But in doing so, she is accepting this picture of subjectivity and objectivity as mutually exclusive. Consider also what Hellman is doing by introducing ideal observers and ideal epistemic conditions: she tells us she is trying to achieve impartiality, and root out bias caused by our limited point of view (e.g. 2011, p. 76). The ideal epistemic conditions and the ideal observer are drawn up as a procedure to transform ‘merely subjective’ emotional and moral responses into acceptable, quasi-objective ones. This procedure mitigates one of the ways in which these judgements are subjective by removing the dependence on a particular point of view. Instead, it transforms these judgements into judgements that could be made ‘from anywhere’, which is as close to ‘from nowhere’ as we’re going to get. And in doing so it also seems to mitigate the second way in which these judgements are subjective: hypothetical neutral observers don’t have sensational responses, and don’t take things to be any way. We should treat all subjective judgements as suspicious until they have been through this procedure. So, Hellman attempts to get to objective knowledge by abstracting away from subjective engagement. And in doing so, she accepts the background view of objectivity and subjectivity as fundamentally opposed.

This underlying metaethical assumption seems to drive the shared assumptions I identified in chapter two. Chambers, Gardner, Moreau and Hellman all believe that while discrimination intuitively involves a personal wrong, a) this intuition as it stands isn’t sufficient to establish that discrimination is distinctively wrong and as such calls for a further grounding, and b) it needs a particular kind of grounding: a reference to a tangible (status) harm. And this underlying metaethical assumption, I will now argue, is problematic.

To start off, this understanding of objectivity isn't the salient one for Hellman's project. Hellman recognises and emphasises that subjective points of view and subjective modes of engagement are crucial to acquire moral knowledge about the wrong of discrimination, and that we should be developing subjective ways of engaging with the world in order to do ethics well. We would be helped, I think, with a metaethics that a) shows us how we can take subjective value judgements seriously, and b) points us toward practical instructions for how to do subjective judgements well. Instead, what we get is a metaethics that fundamentally distrusts subjectivity as biased and not in touch with reality, whose only practical recommendation is to abstract as far away from subjective points of view as possible. Hellman wants to take subjective judgements seriously, but on this metaethics they are the kind of thing that cannot be taken seriously. This is an awkward choice of metaethics that doesn't seem to move her view forward on either a) or b), but instead is a source of doubt about whether we can take interpretative value judgements seriously.

The problem, then, is that this metaethics creates doubt and subsequently doesn't offer the right tools to overcome that doubt. Consider how Hellman's ideally neutral observer would actually go about evaluating a racist policy. As per Hellman's instructions, the observer needs to be both empathetic and impartial. Empathy requires her to assess the situation and take in her subjective response to it. And impartiality requires her to suspend her judgement until she has adequately assessed her emotional response. But how is she to do this? Once suspended, it is unclear based on what further facts or criteria she can draw the conclusion that the judgement is valid. Anytime she would feel compelled to interpret the situation as racist, she would have to remind herself that these subjective responses are parochial and thus can't be fully trusted. Her subjective responses are at most inconclusive evidence that this policy is in fact racist. This seems to leave her no secure ground to make a conclusive judgement beyond 'this is how things *seem* to me, but I might be mistaken'.¹⁵

This metaethics recommends that we navigate the first-personal, subjective perspective by stepping out of it for an 'independent' and objective assessment. But this, I think, is the wrong direction in which to look for practical tools or criteria for validating such judgements. The most accurate subjective judgements aren't necessarily the least subjective ones. The underlying metaethics makes it seem as if we must depend on our subjective ways of engaging with the world only because we are limited, parochial creatures. This is a mistake. These ways of engaging with the world are our only way of determining whether a certain policy demeans not because we are

¹⁵ It almost looks as if this metaethics is gaslighting her: she *knows* something is wrong but she is being made to question her own judgement.

limited and parochial, but because this is the only way in which such facts can be known in the first place. Our being parochial and our need for judgement and empathy isn't a limitation. Instead, these subjective modes of thought enable us to access moral knowledge that isn't neutrally available. They are a strength and a skill, not a limitation.¹⁶ Finding out whether a subjective judgement can be relied upon is not a matter of making them less subjective, but instead of diving into them and examining them more closely.

I've argued that Hellman adopts a metaethical view on which subjectivity and objectivity are mutually exclusive, and that this view doesn't seem well-equipped to either a) show us how we can take subjective value judgements seriously, nor b) point us toward practical instructions for how to make valid subjective judgements. Instead, it seems to place the one casting the interpretative value judgement in a weaker place than where she might've been before we began to look at the metaethical foundations. Hellman and Sangiovanni get something right: we need to say more about the validity of interpretative value judgements, and about which ones are valid. But it seems that we need a different metaethical picture to help us do so. In the rest of the chapter I develop a metaethics on which we can take interpretative value judgements seriously in their own right.

3.3 Wittgensteinian Metaethics

I now turn to a metaethical approach inspired by Cora Diamond, Alice Crary and Raimond Gaita, all of whom draw on Wittgenstein. These authors have more faith in subjective ways of getting at moral truths, based in Wittgensteinian lines of thought, and thus are a good starting place for a metaethics that can underwrite interpretative value judgements. By labelling these views 'Wittgensteinian', I mean roughly this. Wittgensteinian thinkers are interested in what it means to do philosophical inquiry well. They view philosophy as a peculiar and potentially pathological exercise. They are particularly critical of foundationalist tendencies: the tendency to keep asking "why", and the desire to find some kind of Archimedean point of view that grounds all our (ethical) judgements. They are also critical of rationalist tendencies: the desire to find a reason to back up every judgement. This obsession with foundations and reasons can sometimes obscure more than it illuminates and can be downright unhelpful, since it seems our judgements about things is in some sense more immanent. These are the anti-foundationalist aspects of Wittgenstein's views

¹⁶ Consider also that immanent and sound judgement is a marker for moral authority (and perhaps even sanity, as is argued by Gaita (2000)). Neutral observers who succeed in suspending their emotional judgement in the face of racism might have already conceded too much to the racist in order for us to be able to take them (morally) seriously.

that these authors draw on. There is also a more positive aspect of his work that they adopt. Wittgensteinian work doesn't shy away from non-rationalist and subjective modes of knowing. It is interested in the conditions for being able to talk in certain ways – in what we already have to presuppose in order to even just talk and ask the questions we're asking. It sometimes talks about these conditions as 'forms of life'.

Speaking more directly to the point of the chapter, these philosophers are interested in questioning the tendency of analytic philosophy to think that a sound judgement needs to be rationalistic, involving a valid argument referring to neutrally available facts. Judgements don't have to and can't be made from nowhere, should depend on our subjective resources, and can be simultaneously value-laden but objective. Moreover, explanation can stop at judgements. This, I think, makes for the best starting point for a metaethics that can take interpretative value judgements seriously.

3.4.1 Non-neutral moral resources and ethical understanding

How can we validate interpretative value judgements about discrimination, if not through rationalistic argument? To answer this question, I will draw on Crary's metaethics. Crary argues that there are things out there in reality which are a) only accessible in non-neutral subjective ways, b) objective and c) involve values: mental states.

On the mainstream view of objectivity outlined in chapter three, facts are objective when they are neutrally available. On this picture of what our cognitive access to the world is like, our subjective ways of engaging with the world obstruct our view of how things really are (2016, p. 44). So, in order to be justified in thinking we have an accurate image of reality, we need to abstract away from everything that has an essential reference to our subjective responses – such as interpretative value judgements. On this view, objectivity necessarily excludes all reference to subjectivity. In section two, we saw this play out in two ways: 1) subjective judgements are relative to a point of view, whereas objective ones are unrestricted by any such points of view, and 2) subjective judgements are sensational responses, whereas objective judgements represent what there is independently of how we take it to be.

Crary calls this the *narrow* view of objectivity, and rejects it.¹⁷ Instead, we should 'take seriously the possibility that all our modes of thought are essentially informed by subjective responses' (Crary, 2016, p. 56). We cannot access the world from an ideally abstract point of view. So we should take

¹⁷ For her rejection, see her treatment of what she calls the 'abstraction requirement' (Crary, 2016, Chapter 2).

seriously the possibility subjective judgements can be objective, too. For Crary, 'subjective' just means 'essentially related to subjectivity'; some quality's being subjective doesn't make its objective status *prima facie* suspicious.

Crary resists a narrow view of objectivity in favour of a broader conception of objectivity. She argues that some subjective (i.e. non-neutral) ways of getting at moral truths can have objective status. We don't have to deny objective status to qualities we can't adequately capture without reference to subjective responses elicited by objects that possess it – e.g. perceptual or affective responses. Some things, she argues, are both objective (in the sense of being public such that any thinker ought to be able to register it) and have an internal reference to subjectivity. If we attempt to distance ourselves from our subjective responses to the world we lose sight of such facts, and thereby lose sight of a rich domain of empirical ethical facts that should inform our moral thought. Instead, we should see that subjective methods can be valid empirical methods for getting at the objective, value-laden world.

More specifically, Crary argues that mental states are a) only accessible in non-neutral subjective ways, b) objective and c) involve values. She begins by arguing that qualities of mind such as pain are necessarily tied to expressive behaviour (2.2.i). While we can conceal or suppress this behaviour, it is wrong to think that we can sever links between the aspect of mind and its behaviour completely. Following Wittgenstein, Crary argues that it is not clear that the idea of 'purely mental pain' is still recognisable as pain, or even as something negative. This fundamental conceptual link between pain and pain behaviour allows us to directly perceive when someone is in pain. Pain is expressive behaviour. So, when we ascribe 'pain' and other mental qualities, we're not just making an educated guess about how people might be feeling. We're sensitive to how things really are. Aspects of mind are objective and observable.

Crary then asks what is involved in our ability to grasp the significance of expressive behaviour and ascribe mental qualities such as 'being in pain'. How do we observe qualities of mind? In order to do so, Crary argues, we need to make a reference to what is important in the lives of a creature of its kind, and what place that mental quality has in it (2.2.ii). This she calls an 'ethically loaded' conception of human and animal lives. Figuring out whether a creature is feeling pain – whether it's ourselves, other human beings or animals – already involves, and even presupposes such an ethically inflected view of the life of that kind of creature.

This takes some explaining. According to Crary, a grasp of what someone is saying involves a necessary reference to a conception of what matters in human life. When we make sense of what a person is saying, we necessarily draw on 'pragmatic sensitivities'. We can find this to be true in

everyday life: in order to understand whether someone is making a joke or a claim to truth, we need a sense of context. That is, we need to have a sense of the rich variety of things speakers can do with words if we're to take speech acts as intelligible acts. And this sense of the kinds of things speakers can meaningfully do is, argues Crary, inseparable from a sense of the kinds of things that give our endeavours a point (2016, p. 76). We thus need an ethically inflected picture, a sense of the kinds of things that are important in the lives of human beings, if we're to ascribe mental states. So, mental characteristics are ethically significant: in ascribing them we necessarily draw on ethical conceptions of human lives. That is, correctly ascribing mental characteristics 'presupposes' and involves knowing what is important in the lives of human beings; that we know what interests they have.

We get such ethical understanding, Crary goes on, by using our non-neutral moral resources. These are our subjective capacities for arriving at a faithful understanding of human lives, such as our moral imagination and affective responses. One particularly important source of moral knowledge is literature that paints a picture of what is important in human life. Literature has a unique capacity to engage and compel the reader by inviting imaginative exploration of perspectives, shaping their sense of what is important and internally changing how they see the world. In particular, Crary argues that works of literature that invite us to explore new ethical attitudes and perspectives can contribute to a world-guided objective understanding of human beings.

Piecing all of this together, Crary argues that our non-neutral moral thought can be sensitive to reality, and that as such it can be objective, authoritative, and capable of critical scrutiny. Mental states b) are objectively there; we can be right (and wrong) in ascribing them. They are also only accessible by a) using non-neutral methods – we need to view people in a light of what matters in their lives, that is, through an ethical lens in order to even be able to ascribe mental states – which is also how c) values come into it: we have to presuppose this ethical conception of what matters to human beings in order to get at these objective facts. And because something matters to a kind of being, in a sort of deep way, in a way that's conducive to flourishing, it can morally compel us to treat them in certain ways.

3.4.2 Crary, Discrimination and Identity-Based Wrongs

On Crary's picture, we need our subjective capacities in order to bring human beings empirically into focus in a way relevant to ethics. Crary's metaphysics are controversial and leave a lot of

questions. I will limit myself to discussing two of them. First, what are we supposed to learn from exercising our subjective capacities? Second, how can we tell when we get it right? I will discuss the first question in this section, and the second in chapter four.

When we exercise our moral imagination and use our affective responses, what empirical knowledge do we gain? Crary sometimes writes as if we learn from reading books that certain ways of treating people are wrong or abhorrent – but while that may be true in some sense, it can't be this straightforward. The metaphysical work she has done does not directly validate moral affective responses like 'this treatment is wrong'. If we grant Crary the metaphysical argument above, we grant that our non-neutral moral resources can get us (and presupposes) ethical understanding: knowledge of mental states and with that knowledge about what is important in the life of human beings. We don't yet know, however, how this ethical understanding translates to the conclusion that treating human beings in certain ways is wrong. This intuition – this treatment is wrong – is supposed to be validated by the argument; it can't be used as a premise.

Instead, what we learn (and already need to know in order to be able to understand human beings) in a general sense is the following: that all human beings, however well-endowed cognitively, figure in moral thought as beings who, simply in virtue of being human, are morally important, meriting specific forms of respect and attention (Crary, 2018, p. 18). We learn through our non-neutral moral resources that human beings are the kinds of creatures to whom a certain treatment is important.

With Crary, we can say that human beings' interest in being treated equally shows from their human psychology. Crary describes how this could work (2016, p. 83):

Suppose that we consider what children are doing when they are first learning to think and talk about love and hatred. Such learning presupposes modes of interest and attention that reflect a sense of the importance of things like reliable companionship and social acceptance.

Crary writes that our talk about love and hatred reflects the importance of social acceptance, and we can make a similar argument about status equality. Many kinds of relational attitudes and emotions that are central to human life also presuppose an interest in status equality, like humiliation, respect, dignity, and pride. These are ideas that we can only understand emphatically by presupposing human beings have an interest in equality. To question whether human beings have an interest in status equality puts us in a position from which we simply can't understand talk of those things. Knowing that human beings have an interest in status equality is *both* a prerequisite

of understanding them *and* something we can learn from paying close attention to them using our subjective capacities for understanding them.

At this point we might balk, as it seems that drawing up a notion of what is important to human beings will always be a circular enterprise: we find out what is important in human life by looking at it through a conception of what is important in it. We already know what we're supposed to learn. But this, argues Crary, needn't be a problem, as she has the theoretical tools to solve this circularity (2016, p. 178): in arguing for allowing ethical understanding into objective thought, she has already argued that we can't bring the world into view in a way that abstracts from our ethically inflected knowledge. Our cognitive access to the world has to start from somewhere, and so being dependent on a non-neutral starting point isn't an argument against her – or any – view. There is no *prima facie* problem with trusting our ethically inflected view to be in touch with reality and to tell us what is important to human beings as a kind. On the contrary, by trying to stand outside of this ethically inflected view we close ourselves off from knowing what human beings are like.

So, through our non-neutral moral resources we discover that human beings are the kind of creature to which equal treatment is important. From the ethical judgement that human beings have certain interests it is a small step to the moral judgement that to treat them otherwise is wrong. This step of the argument turns on accepting that human beings merit the things that are important to them. Why would this be so? Merely finding something important doesn't seem sufficient to elicit strong moral obligations: I can find it extremely important to be addressed as Great Leader, but that doesn't oblige anyone to actually treat me in this way. Instead, Crary is aiming for a specific moral way in which things can be important to us: the things that are part of what it is to live a good life as a human being. What things are important in this moral sense depends on the kind of being that you are. In the case of human beings, the argument goes, we learn from our moral resources that being treated in certain ways is part of what it is to live a good life. Human beings, then, merit to be treated equally because that treatment is morally important to them, that is, something that make their lives go well.

Note that even someone who does not or cannot register this equal treatment (or a violation of it) merits to be treated equally all the same. The argument doesn't depend on the person being in any tangible way affected by the treatment they receive. Equal treatment by itself just is what makes people's lives go well in virtue of being human beings, and unequal treatment by itself makes their lives go worse.

This explains what goes wrong if we refuse to look at human beings in an ethically inflected way and demand a narrowly objective feature of the world to substantiate our intuitions. In questioning

whether human beings have an interest in being thought of as equals, you are effectively side lining yourself from talking in a meaningful way about human beings. We have to presuppose this interest if we're to understand human beings at all, we have to be in that circular way of looking at human beings. The fact that human beings have an interest in being thought of as equals isn't the kind of thing that requires justification.

It also gives us a positive argument for the moral importance of treating human beings with equality. Changing the metaethical approach allows us to rephrase what kind of story would allow us to accept this intuition as morally informative. Instead of asking 'Why are identity-based harms distinctively wrong?', I can now ask: Does the intuition that identity-based harms are distinctively wrong tell us something about how the world really is, or is it just a misleading sentimental response? I've suggested that it does tell us something about what the world is really like. We are the kind of beings that have an interest in being treated equally that this intuition catches onto.

Using our non-neutral moral resources and moral imagination, we can make objectively valid moral judgements. These judgements tell us something about and draw on our knowledge of the kind of beings that we are: the kind that has an interest in being treated equally. These judgements can function like axioms that don't require further justification to accept them. They are, like mathematical axioms, the justificatory bedrock with reference to which all other propositions are or can be justified. As such they form the basis of a theory of wrongful discrimination, and more broadly, an account of whether identity-based harms and wrongs are distinctively wrong.

3.4 Conclusion

In this chapter I set out to defend the status wrong view of the wrong of discrimination and identity-based harms and wrongs more generally. With Hellman, I argued that we can have true judgements about when people are being treated as inferior in a way that offends against equality. I argued that this view is detailed enough to help us distinguish discrimination from morally permissible kinds of treatment (3.1). However, I argued that Hellman's account of modest objectivity fails, and with it her account of why we should take these judgements seriously (3.2). It fails because even though it recognises the need for subjective ways of knowing, it still subscribes to a picture of objectivity as a 'view from nowhere'.

I then put forward a Wittgensteinian metaethics on which subjective resources are necessary for ethics; not *prima facie* distorting. I challenged rationalist and foundationalist assumptions present in the debate on status wrongs (3.3). With Crary, I put forward a metaethics on which subjective

interpretative value judgements can have full objective status (3.4). Changing the metaethical approach allowed me to rephrase what kind of story would allow us to accept this intuition as morally informative. Instead of asking ‘Why are identity-based harms distinctively wrong?’, I can now ask: Does the intuition that identity-based harms are distinctively wrong tell us something about how the world really is, or is it just a misleading sentimental response? I suggested that it tells us something about how the world really is.

Chapter four – Getting it right

I have put forward subjective resources as necessary for ethics; not *prima facie* distorting. And I have suggested that the intuitions surrounding the personal wrong of discrimination capture what the world is really like. However, any particular conception of what matters in human life may turn out to be distorting. So how can we tell the difference between a sentimental or distorting conception or way of bringing someone into focus, and one that gets it right?

I will explore this question by engaging with the racist conception of Black, brown and white people found in Jean Raspail's *The Camp of the Saints* (1973), a canonical novel in white nationalist circles. On the face of it, Raspail uses the techniques Crary praises for being able to reveal aspects of the world otherwise hidden from view. His novel aims to elicit certain affective responses, foster kinship, and expand the reader's moral imagination beyond what is normally accessible. It is an attempt to advance particular ways of looking at human beings by bringing their lives empirically into focus. His message, however, is one we need to be able to reject. Raspail's novel is deeply racist: it fosters kinship between white people, and invites white people to respond to people of colour with disgust and fear.

On what basis could we reject Raspail's work? One standard response to racist views is that they play on our emotions instead of on our reason. For example, we might say that by creating an environment of fear and anger, racist politicians manage to play on people's emotions rather than the actual facts in front of them. But this response is unavailable to us. I've argued with Crary that sentimental responses are legitimate, objective ways of bringing the world empirically into focus. A literary description that invokes emotions *is* a rational argument. So merely saying that racism plays on people's emotions instead of the facts or our capacity for reason is a false dichotomy; they are, or so I have argued, one and the same.

What we want to say is that the emotional responses Raspail elicits are distorting, misleading, and false. But so far, it's not clear how we could do so. In embracing subjectivity, it might seem, we have lost the objective standards or criteria we could otherwise appeal to in deciding what the world is really like. As long as Raspail can convince people to have certain emotional and sentimental responses to reality, the thought goes, we should treat his perspective as at least potentially able to reveal otherwise hidden aspects of the world. He is speaking 'our language', so we need take him seriously.¹⁸ If this worry isn't met, we have ended up with a metaethics on which

¹⁸ But *is* he speaking our language and are we even able to take him seriously? With Crary I've argued that we need to presuppose an interest in equality if we're to talk in a meaningful sense about human beings. The racist isn't doing

anything goes as long as you can foster the right sentimental reactions in people. And that cannot be right.

There is a further, more theoretical worry here. With Crary, I'm claiming to have introduced a new kind of ethical argument – the appeal to our non-neutral resources. However, as this challenge makes clear, I have not yet provided any argumentative rules. Without such rules as to what counts as a good argument (and what counts as a bad one), we might wonder whether we can really say we've added a new kind of ethical argument to the mix at all. If there are no rules of application, no ways of deciding which contribution is valid and which contribution is wrong, it seems that there is no argument in the first place.

This is the challenge. My aim in this chapter is twofold. First, I investigate how literature that invites racist views of human beings might differ in relevant ways from the kinds of pictures of human life I suggest are authoritative. Second, I develop Crary's non-neutral metaethics by specifying how we can tell veridical sentimental responses from misleading ones.

The description I will give from *The Camp of the Saints* is viscerally racist. Before I start, I want to explain why I am engaging with it. Perhaps racist views don't merit serious discussion. Perhaps considering racist views too seriously should make me fear I've lost my capacity for proper judgement or my moral authority.¹⁹ And perhaps discussing racist ideas does more harm than good, for example, by promoting this racist view as a serious alternative, spreading racist ideas, and constituting racist trauma voyeurism. What good can there lie in taking the racist seriously?

First, Raspail's narrative is powerfully compelling to a growing number of Europeans and North Americans, as far-right and alt-right white nationalist ideas are becoming increasingly mainstream. Raspail is heralded by white nationalists as one of the first to take seriously the threat of mass-migration to Western ways of life. Given its popularity, it seems important to understand this ideology and know what we want to say to and about those espousing it. And I think that my preferred metaethics can do a better job than what we have available now. Many people respond to visceral racism with a visceral affective response: that's just a disgusting, inhumane way of looking at people. On mainstream metaethics, we're required to back this response up with facts. Such as: look, science tells us people aren't different, Black people suffer too, et cetera. But it

this. So presumably we could just dismiss racist views as nonsense. Gaita takes this idea one step further by arguing that some people have 'so radically lost his capacity for judgement that his views aren't worth considering' (Gaita, 2000, p. 3). The racist might seem like he is speaking our language on the surface, but he is speaking in a language on which words like 'argument' and 'evidence' have lost all meaning. If this is right, my point should merely be that it *seems* that that the racist is speaking our language.

¹⁹ See also my discussion of the ideal judge in 3.2.2.

seems to me that these facts don't fully capture our moral response to racism. This is not what we really want to say. We want to say that the racist view is disgusting; that we are talking about *human beings* here. This more emotional response is the one I aim to validate. Second, my preferred metaethics makes it so that I need to take this racist ideology seriously, because on the surface its argumentative form *looks* a lot like the kind of argument I am endorsing. If I am promoting affective responses as valid ethical arguments, I'd better have a story about why the racist ones don't count. And third, I need a story about the argumentative rules in order to make this metaethics plausible. In that sense, my discussion of the racist ideology is a proxy discussion for the whole method.

In section 4.2, I lay out how *The Camp of the Saints* invites us to look at white and non-white human beings. In section 4.3, I examine Crary's discussion of what makes a picture authoritative and discuss the problem of conservatism. I suggest that we can reject racist literature such as *The Camp of the Saints* because it refuses to pay attention to the lives of Black and brown people in light of what matters to them. In section 4.5 I discuss the positive-bias problem: more attention doesn't always seem better. I conclude that some forms of attention are ethically frivolous.

4.1 Reading *The Camp of the Saints*

4.1.1 *The Camp of the Saints*

The Camp of the Saints tells the story of the destruction of white Western civilization through mass immigration. In the book, one million Indians assemble at the Ganges, board any ship they can find and slowly but surely make their way to France. I will outline how Raspail means to change the way readers look at human beings of different races by appealing to their sentimental responses and opening up their moral imagination to new ways of looking at human beings.

First, Raspail invokes images of disease, disability, excrement, stench, sex, death and liquid to evoke disgust for the Indians on the fleet and dehumanize them. He invites his white readers to look at these Indians as poor, wretched, deformed, disabled and diseased, describing the children as 'wretched creatures'; 'the damaged ones, starting to rot, all wormy inside, or turned so you can't see the mold' (ch. 5). The spiritual leader of the fleet is the disabled 'monster child', and (physical) disability is widespread among the Indians. The Indians are also associated with excrement and stench. The father of 'the monster child' is a dung roller by trade only referred to as 'the turd eater', and India is described as smelling of burned excrement (ch. 20). The fleet trails the stench of

latrines and corpses: 'the whole ocean was like one big festering sore', 'a welter of dung' (ch. 27, 20). And even those of Indian descent living in the UK, who are presumably more Westernised, have a 'rather pungent and unfamiliar smell' (ch. 35).

Bodily liquids – like puke, sweat, urine, drool, feces and sperm – also serve to evoke disgust, and serve to further dehumanize these people (ch. 43):

As the decks sprang to life with their myriad bodies – men, women, children, steeping in dung and debris since Calcutta – as the hatchways puked out into the sunlight the sweating, starving mass, stewing in urine and noxious gasses deep in the bowels of the ships, the stench became so thick you could practically see it.

As the masses are dying they become even more liquid: they get more slimy and wet, their flesh starting to melt and drip (ch. 20):

Endless piles of burning bodies – dank, slimy entrails, mostly – overspreading the sea with a noxious stench. Arms and legs tumbling from the pyres, too narrow for their loads. Heads of scorched hair, rolling at the feet of the squatting multitude ... The cremators would poke long boathooks into the mass of molten, dripping flesh, and try to push it neatly back in place.

Second, metaphors of flow and waves serve to simultaneously dehumanize the Indians and underline their threat to the West. The Indians are continuously described as indiscernible mass of moving, flowing, cascading bodies, a horde, with no distinction between one refugee and the next, a swarm, 'an endless cascade of human flesh'. The mass moves like a liquid: they are an unstoppable wave, an out-of-control and flood of human flesh and bone, overflowing from the Third World with the West as its sewer. Images of flowing body parts spill over into imagery of group sex, which has the additional function of painting these people as savages or animals. And once the mass has started flowing toward the West, the flow is uncontrollable, with the human cascade pouring down the sides of the ships, 'swelling into huge wave upon wave of flesh, bodies upon bodies, pushing, shoving toward the shore', a 'flood of flesh and bone' (ch. 43, 27).

Third, Raspail uses distance to signal white people's blindness to the truth. Distance allows Westerners to look at Indians with compassion: they use the language of compassion and anti-racism, adopt babies and send humanitarian help. With a vast ocean still between them, the West is able to portray the fleet as pitiful and as a 'splendid example of human dignity' (ch. 21). As this distance closes in, however, the white characters are increasingly less successful at masking what Raspail suggests is their true and proper emotional response to the Indians: shock, terror, disgust

and a ‘good, healthy fear’ (ch. 22). And as the ship comes closer to French shore (as we take a closer look at these human beings), the Western world starts to panic and flee. Language might be powerful and able to distort how we see these human beings, Raspail suggests, but if we look up close our true affective responses win out.

All this stands in stark contrast with Raspail’s description of Westerners. First, Raspail paints a picture of white Western kinship and culture. France, *the camp of the saints*, is a place of prosperity and abundance, of culture, elegance, refinement, and civilization. Perhaps the best symbol of this culture is the first character we meet, the old professor Calgues, an intellectual steeped in history and culture. Calgues feels joy, happiness, pride and self-respect as, Raspail suggests, whites can. He is confident in his white Western superiority: ‘the knowledge that one’s own is best, the triumphant joy at feeling oneself to be part of humanity’s finest’ (ch. 1). And he thinks of kinship as racialised (ch. 1):

‘Man has never really loved humanity all of a piece – all its races, its peoples, its religions – but only those creatures he feels are his kin, a part of his clan’.

Second, Raspail means to show his white readers that this strong, powerful and dominant race is deeply vulnerable to threats from outside and within. From outside comes the mass influx of migrants who, as Raspail stresses, greatly outnumber white Westerners. Even if this particular fanciful fleet hadn’t set for France, Raspail implores, ‘with millions of us and billions of them, we couldn’t have held out much longer’ (ch. 40). This white vulnerability is compounded by a political atmosphere that has weakened them from the inside: the rhetoric of non-discrimination, human kinship, unity and compassion. As such all things important in the lives of white Westerners – strength, pride, history, glory, wealth, a sense of achievement, a place of one’s own, a sense of security and identity and superiority – are threatened by migration. The Western race as a whole, Raspail suggests, will be assimilated and cease to exist.

4.1.2 *Different Kinds of People*

It is helpful to compare how Raspail speaks of different kinds of human beings with how Cora Diamond speaks about different kinds of animals in *Eating Meat and Eating People* (Diamond, 2016). Diamond, like Crary, argues that we cannot do ethics with a neutral view of human beings and animals. Instead, we need to find reasons for treating them in certain ways that are infused with human imagination and sympathy. We can find such reasons if we look at animals as our fellow creatures whose company we could seek out, and who we could pity.

As both Crary and Diamond point out, we have a wide range of feelings and behaviours directed at different kinds of animals. Crary notes that since we keep dogs as pets, companionship relations with dogs are familiar to us (2016, p. 260). For Crary, the close relationship means that it's easier for us to access facts about the lives of dogs. This makes attending to dogs a good starting point for developing an ethical conception of the lives of all animals, such as, for instance, animals killed in factory farming. But that relationship doesn't itself seem to influence how we should be treating dogs. What treatment an animal merits isn't conditional on the animal's relationship to us, but on what a being of its kind needs to flourish.

Diamond, however, seems more deeply invested in viewing the particular ways in which we relate to different kinds of animals as a starting point for animal ethics. From this starting point it makes sense to treat animals differently based on our relation to them. This 'relating to' is the only possible starting point for ethics, and we need to start from this subjective way of relating to animals if we're to make any convincing ethical argument about how to treat them at all. In ethics, we have no choice but to start with this given plethora of relationships. Note that saying that we relate to animals in different ways is not to say that these kinds of animals are natural kinds; it just means that we cannot relate to these animals without thinking of them as pets, prey, or vermin and ourselves as loving owners, fair hunters, and threatened exterminators. Animals always show up for us *as* related to us in a certain way, and that this deeply subjective 'relating to' is our starting point for ethics.

According to Diamond, we relate to most animals as fellow creatures in some way. But viewing animals as our fellow creatures also sometimes goes hand in hand with killing and eating them. One example Diamond gives is the idea of a fair hunt in which the hunted animal is viewed as a respected enemy (Diamond, 2016, p. 102). Another is the idea that it's okay to eat farm animals as long as they are treated respectfully and not hurt 'more than necessary' throughout their lives (Diamond, 2016, p. 96). Even though you kill and eat them, animals are seen as fellow creatures: they have a life and are individuals of whom it makes sense to say that you could spare them out of pity. But still, our treatment of them differs widely depending on how they feature in our thought. These complex and differing relationships to different kinds of animals are morally important.²⁰

²⁰ It's ambiguous, reading Diamond's *Eating Meat and Eating People*, whether we ought to spend energy to extend our view of animals as fellow creatures to new animals, extend our view of animals as not to be killed to new kinds of animals. Diamond tentatively suggests that 'imaginatively reading into animals something like an appeal to our pity' can help us grasp their situation more fully (Diamond, 2016, p. 107). Doing so is to frame their lives in terms of a sense of having importance and loss. But there is also a conservatism at play here: this appeal to pity only works if an animal is in the range of our fellow-creature response.

The analogous suggestion we might take from Raspail is that Crary's focus on fellow human beings obscures some other fellowship relations that are also morally salient. Crary wants the typology of 'human being' and 'animal' as a kind to do the heavy moral work. People show up for us as human beings, and we should view them through that ethical lens. Similarly, animals show up as our fellow creatures, and (argues Crary) we should view them through that ethical lens and refrain from killing and eating them. But, Raspail argues with his novel, there are significant forms of kinship within the group of human beings too. In our heavily racialised society people show up as Black, brown or white, and we (sometimes) relate to people as such. And as we do, our responses to them differ too, particularly, Raspail suggests, those of kinship and fear and disgust. White people, Raspail argues, feel a special kinship towards other white people. Because of political correctness, these feelings are often hidden and distorted. But, Raspail tries to show his white readers, the feelings are really truly there: if they look close enough, they do feel that joy of being part of a superior race, with a superior cultural heritage. And similarly, Raspail argues, once white people remove the distance between them and Black and brown immigrants – as Raspail dramatizes by letting the fleet come closer and closer to France – their responses of fear and disgust show that white people relate to Black people very differently than they do to other white people. Merely thinking of morality as how we relate to our fellow human beings, we can read into Raspail, is to have the wrong moral typology. The typology of the kinds of beings that we need to appeal to in order to sensitize to the relevant moral facts isn't (always) species, but (sometimes also) particular social identities, such as race.

Just like with Diamond's discussion of different kinds of animals, saying that race matters morally is not necessarily to say that races are natural kinds.²¹ It just means that we cannot relate to people without thinking of them as standing in a particular racial relation to us. (Or in any case, that if we fail to so in ethics we're missing out on something crucial). In this case, it means that we cannot do ethics without thinking of people and our kinship to them as racialised.

Crary argues that ordinary philosophers miss out on ethical significance by refusing to look at human beings and animals in a light of what matters in their lives: you need this subjective perspective in order to discern the ethical significance otherwise hidden from view. Perhaps the same is true of being a member of particular social groups. Most ethics is done from a generalist human perspective – 'we're all human' – suggesting that we're all subject to the same ethics of being human beings. But perhaps, the suggestion is, this generalist point of view disables us from viewing ethical significance of social identities that is truly there.

²¹ Though Crary might argue 'human being' is a natural kind.

The suggestion is that since our relations to other people are racialised, we should take this racialised typology as a starting point for doing ethics. This idea is familiar to anti-racist activists, too. The activists I have in mind strongly oppose the idea of colour blindness – the idea that we should stop seeing colour and all just relate to each other as human beings (see e.g. Baldwin, 1962, pp. 405–408; Eddo-Lodge, 2014, p. 83). They reject the idea that to solve the problems of racism we should just stop relating to people in a racialised way. Race matters to our personal interactions and identities, and racism structures the lives of Black, brown and white people. Giving a ‘neutral’ description of people’s lives as devoid of race completely obscures real aspects of life.

Raspail, then, can be read as putting forward the idea of racialised kinship between white people. According to Crary ‘our fellow human beings’ is a morally salient category. Raspail’s suggestion is that ‘our fellow white people’ is so too for white people, and that white people relate to other white people differently (with pride) than to Black and brown people (with fear and disgust). Overlooking these facts, Raspail argues, is to miss something morally important.

4.2 Authoritative Perspectives

4.2.1 Crary

Central to Crary’s account is the power of literature to bring the lives of human beings ethically into focus. However, as she recognises, skilfully eliciting certain responses in people isn’t enough to make for the kind of argument we’re after, as any particular response might be distorted. We should always take a step back from the ethical perspective presented to us and critically assess it. Valid perspectives are those ‘capable of surviving critical scrutiny and establishing themselves as authoritative’ (2016, p. 82). The main question for Crary is whether the perspective is a mere projection, or whether it enables us to see things that are really there (e.g. 2016, pp. 138, 144, 148, 162). But how can we tell? Here are three core features of authoritative perspectives.

First, authoritative perspectives enable us to *make sense of otherwise puzzling observations*. Crary has in mind specifically observations about how we treat people, and how we respond to their maltreatment in return.

I want to highlight one such case that Crary considers: the case of how we treat people with intellectual disabilities. In the face of philosophers who hold that people with intellectual impairments matter less morally in virtue of those impairments, Crary describes how people often

respond with particular indignation to the abuse of people with intellectual disabilities. Borrowing from Cora Diamond, Crary writes (2016, p. 131; Diamond, 1991, p. 56):

Diamond asks us, for instance, to consider the kind of indignation “we may feel at the rape of a girl lacking speech and understanding, lacking what we think of as moral personality and the capacity for autonomous choice, and incapable of finding the event humiliating and the memory painful as a normal woman might.”

This imaginative exercise is meant to illustrate that we often respond as if the abuse of those who aren’t able to conceptualise that abuse is – if anything – *worse* than the abuse of someone who is cognitively able to. The fact that this girl cannot understand the rape and might not experience it as a traumatizing attack doesn’t mitigate the wrong of rape. The rape still violates her bodily integrity, her need for safety, her sexuality, her independence as a person, and her need for companionship. Moreover, her inability to conceptualise the rape as rape might exacerbate these wrongs. This girl is someone who is unable to enjoy some things we think makes human life good. And she is less able to assert or show that these things are important to her, and thus more dependent on others to provide her with the companionship, respect and attention she is due. This ‘special outrage’ we feel at her abuse, writes Crary, can only be made sense of if we adopt an ethical perspective on which people with intellectual disabilities are fully human, meriting companionship and safety.

Second, authoritative perspectives allow us to *see more*. They reveal significant aspects of someone’s life that are otherwise hidden from view. By adopting the right ethical perspective ‘we position ourselves to discern things that aren’t otherwise available’ (2016, p. 136). Without this perspective we could be ‘missing significant features of human life that are truly there’ (2016, p. 139).

One example of features hidden from view is given by Crary’s discussion of the lives of the seriously cognitively disabled, such as Fred Brandenburg.²² Brandenburg lived in a private home where he was structurally neglected and abused, eventually leading to his death (2018, p. 6).

Having failed for years to give him recommended medication for a heart condition, staff at his home tranquilized him without any medical indication one day in 1997, when he was 57. The effects on his health were dramatic and dire—within days he was sweating and shaking, unable to stand or eat without assistance—but no one called 911. When paramedics were finally summoned, he had already been dead for hours.

²² <https://www.washingtonpost.com/archive/politics/1999/12/05/invisible-deaths-the-fatal-neglect-of-dcs-retarded/1f820d00-60f8-4f42-9805-b27a84178b86/>

Brandenburg merited medical treatment and freedom from abuse. But it is only by looking at Brandenburg's life from the right ethical perspective, Crary suggests, that we can see the extent of the callous way in which he was treated. By paying attention to the fact that Brandenburg has a mother who named him, and by describing how he spent his hours printing his name on old telephone directories and TV Guides 'impress[ing] the fact of his existence on his world', Crary invites us to see that Brandenburg was a fully-fledged human being who merited companionship, dignity, and respect. He had to trust staff to take care of his health, but also of these basic human needs. Instead, he was mocked and abused. This violation of his trust and his dignity – one that he probably wouldn't even have been able to conceptualise – can only be fully seen by adopting the right ethical perspective.

We have now seen two ways in which authoritative perspectives advance our understanding: making sense of otherwise puzzling observations, and allowing us to see more morally salient facts. These aspects of understanding build on each other. We can explain our responses by 'seeing more', by being increasingly sensitive to the situation in virtue of adopting a better ethical perspective. And if we see more and get increasingly sensitive, we're able to finetune our affective responses to particular situations.

It seems that there are two subtly different ways in which authoritative perspectives help us to make sense of puzzling observations and see more. For one, such a perspective may take us just beyond the bounds of our ethical thought. It might, for example, help us see that mocking cognitively disabled people is wrong even when these people aren't sensitive to this wrong themselves. This vindicates the responses of people who have already caught on, like many parents with cognitively disabled children. But Crary also gives us the tools to understand things which are just completely puzzling, and of which we might ask why anyone would respond in this way. For example, we might accept the idea that people with cognitive disabilities merit care, but not be able to understand why certain things *count* as caring in the first place. Even someone who cares a lot about a person with a severe late-life cognitive disability might not understand why anyone might find it important to dress them in a way that they would've deemed respectable before they had their disability. There is some fascinating way in which this is a senseless gesture, and it might look utterly strange, excessive or overly sentimental. But at the same time, such gestures can look incredibly ethically compelling. The philosopher's urge is to take ethics in a direction that carves this stuff away, and focus on systematic and rationalised judgement that moves us away from these weird ideocratic responses. But something in us wants to take these rich responses seriously and see if we can understand and vindicate them. Crary's and Diamond's metaethics gives us the tool to do so.

This brings me to the third feature of authoritative perspectives: our increased sensitivity and ability to make sense of our responses must come hand in hand with emotional development. Paying close attention to someone's life can shape our affective responses only if we're receptive. Crary makes the argument for this by way of discussing Tolstoy, Coetzee and Sebald's novels in which characters with stunted emotional growth are unable to understand themselves, and unable to respond properly to others. Even if these characters try to 'see better' and make sense of how they or the people around them respond to the world, they do not succeed until they go through a period of emotional growth and develop a certain emotional maturity. For Crary, novels make rational arguments. This means that our moral capacities, such as our moral imagination and affective endowments/emotional responsiveness, are intellectual skills that need to be properly developed for us to be able to bring the world empirically into focus. Not realising this puts us in danger of stunted moral growth, and a stunted view of the world as a consequence.

Crary gives us some examples of stunted emotional development. First, a traumatic break from one's past (as with Austerlitz, 2016, sec. 6.4). Second, a focus on scholarly learnedness rather than cultivation of moral sensibilities (as with Tolstoy and Coetzee, 2016, secs. 6.2, 6.3). Third, having 'social privileges [that] are marks of complicity with society's injustices' (as in Coetzee's work, 2016, sec. 6.3).

4.2.2 Seeing more

According to Crary, authoritative perspectives 1) enable us to make sense of practices and behaviour, 2) enable us to see more and 3) can only be grasped if our emotional sensibilities are well-developed. How does Raspail's racist view of non-white and people hold up under these criteria?

First, authoritative perspectives enable us to account for people's responses to particular situations. It's not hard to see that a racist ethical perspective, too, enables us to make sense of widespread practices and behaviour. Raspail's perspective enables us to make sense of the backlash against immigration, a hatred for Black and brown people, and the celebration of white kinship by increasingly large groups of white people.

This brings out a deeper problem for Crary that I will call conservatism. Crary's first criterion is that authoritative ethical perspectives allow us to make sense of the ways in which we respond to people. But what is to say that those ways are morally sound and cued into reality? People respond to each other in all sorts of ways that don't reflect any aspect of reality, some of which are morally

highly problematic. The conservatism worry is that in adopting this criterion, we run the risk of merely legitimizing people's current responses without critical appraisal. If all we do is explain people's current responses, it seems we lack the tools for progressive moral thought and uncovering moral facts.

Second, how about *seeing more*? Raspail's racist views can be thought to enable us to *see more*: as discussed above, Raspail's perspective allows us to see the ethical significance of race and racialised interactions. More specifically, Raspail is heralded by his readers as making visible a threat to white identity that many can't see. Even the structure of the novel is set up in this way, by making the migrants move closer to France with white people's responses getting more 'accurate' as they come closer into vision.

At first brush, "seeing more" as a criterion doesn't seem helpful to distinguish veridical from misleading racist ethical perspectives either. There is, however, another aspect to Crary's assertion that such perspectives allow us to see more that I believe will be more helpful in understanding why racist views are wrong. I have described how an ethical perspective literally enables us to see things that we would otherwise miss out on: we need the conceptual apparatus to be able to see the moral importance of dignity in Fred Brandenburg's life. But literature that cultivates a particular perspective also allows us to *see more* in a more banal sense, simply by inviting us to pay closer attention to the lives of human beings. Simply by spending time and paying close attention where dominant culture suggests there is nothing to see, we can get someone's life clearer into focus. 'Seeing more' is an invitation to stop and look at what is in front of us.

It is this sincere attention to the lives of the human beings it claims to talk about that Raspail's novel utterly fails on. *The Camp of the Saints* works by obscuring and dehumanizing the lives of one million migrants, and by extension billions of people. The migrants aren't individuals. None of them, for example, have names. The disabled child and his father are described as individual characters, but they are referred to as 'the turd eater' and 'the monster child'. The same happens in the storytelling. The reader doesn't get invited to pick up the Indian's perspective, with all of the story being told through a white storyteller and white characters. All this stands in stark contrast to the white characters we encounter in the book, of whom we learn what they are thinking, and whom we follow as individuals traveling across the country, making life decisions, and talking to their loved ones and enemies.

The suggestion is that we don't need to take the perspective espoused in *The Camp of the Saints* seriously because it is actually a terrible example of the kind of book that Crary has in mind. It doesn't illuminate the lives of these migrants, nor give us any sense about what is important to

them. It creates white kinship by obscuring and hiding from view the humanity of Black and brown people. It doesn't seriously pay attention to what the lives of these human beings is like in light of what is important to human beings.

Raspail might want to suggest that white people might see more through the othering of non-white people: white people can see 'them' for who they 'really' are, which in turn enables them to experience white kinship. But this seems wrong. Seeing human beings necessarily means seeing them in a way that's somehow individuated. If you see people as essentially a mass or a mob, you fail to see them in a way that brings into focus their most morally salient features. You have to see them as individuals on some level, or else you're at risk of seeing them in a distorted way. Moreover, othering is never a good example of paying due attention to someone. It is, on the contrary, the main technique used for those who aim to overlook certain people.

Racist literature shouldn't inform our ethical perspective because it makes no serious attempt at illuminating the views it espouses. Good literature, for our purposes, is detailed and looks closely and attentively at the lives of human beings. We can reject novels that refuse to do this and instead work to obscure aspects of people's lives.

Finally, Crary suggests that having social privileges and being complicit with social injustice can lead to impaired emotional development, which in turn can bar people from seeing the world as it is. This is likely the case for Raspail's characters, and so could function as a defeater for the views they espouse. However, we do not seem to have enough information to make this judgement: there is no in-depth character assessment of any of the characters, and we can't judge some character's emotional development based on the bad view they hold and then use that judgement as further evidence that the view they hold is bad. In addition, we would need well-developed account of healthy emotional development that deals with the complexity of our emotional life, and accounts for the fact that different emotional sensibilities can serve different goals. Such an argument, if it could be made, would require a book on the psychology of racists. I will not make it here.

4.3 A problem with attention

4.3.1 *How much attention is too much?*

I have suggested that a useful criterion to separate qualitatively good literature from the bad literature, and thereby also separate helpful perspectives from racist ones, is whether proper attention is paid. In this section, I discuss a problem with this view.

Understood in one way, Crary's account leans too heavily toward a conservatism: it runs a risk of explaining how people respond, instead of discerning how we *should* be responding. I suggested that literature that pays more attention to the lives of individuals provides qualitatively better arguments, and should win out over accounts that refuse to do so. But this account faces a new problem: it seems to have a positive bias. It favours ethical perspectives that say 'there is something here', as opposed to, 'there is nothing here'. It makes it easy to argue that there is something we've overlooked and should be paying attention to, and incredibly hard to argue against that position – to say that someone mistakenly sees significance where there might not be any. There is a strong bias toward revealing things, even if those things might not be there.

Crary seems to support this inclusive and positive-bias way of thinking that finds significance anywhere. She writes that we have 'overwhelming moral reasons' to refuse to think that creatures aren't imbued with ethical significance, unless 'in some awful and earthshaking case, we find that there is really nothing else to think' (2016, p. 164). As long as there is something to think, we should be thinking it.

I'd like to flesh out this challenge by looking at Koechlin's views on plants (Koechlin, 2009).²³ Noting that 'anything and everything can be done with plants today', Koechlin sets out to broaden our moral horizon. Plants, she argues, are not objects automatons (it) but sensitive living beings (she); they sweat, remember, learn, engage in lively relationships with their peers and environment, and when she walks among them she hears 'a whispering and murmuring of fragrance'. As such, it makes sense to think of them as having dignity. This, she argues, means they merit not to be harmed in an arbitrary way. As examples of impermissible behaviour toward plants, she floats ideas like the 'senseless picking of a roadside dandelion', 'the massive and total instrumentalization and industrialisation of plants' and giving them control over their own genetic heritage. In short,

²³ <https://www.youtube.com/watch?v=i8YnvMpcrVI>

Koechlin suggests that based on the kind of sentimental argument we've been pushing, plants merit certain kinds of attention in a morally deep way.

Koechlin says that attributing dignity to plants 'of course' doesn't mean that we should not eat, cut grow, or mow plants, pointing out that attributing dignity to animals didn't mean that we took them out of the food chain, or quit doing animal research. But this is not obvious at all. Many, like Crary, argue strongly that we shouldn't treat animals in these ways, and it seems Koechlin is setting us up to make similar arguments about plants, noting that we might in a few years laugh at our own arrogance with respect to them.

The problem this poses is this. How do we know when to look for more, and when to stop looking? There are two worries here. First, the worry is that we're inventing things, with our sentimental responses in overdrive, (moral) imagining away and inventing significance where there isn't any. We are just imagining plants' 'whispers and murmurings'. A well-delivered story might make us feel very strongly about the moral fate of plants, but we might also still think that this sentiment is misleading. Second, more subtly, the worry is that focusing on these details leads to the exclusion of other things. In paying too much attention to plants, we might get thinking about establishing moral rights for plants that impact significantly on human welfare. This attention seems ethically frivolous, in the sense that we might be so obsessed with them that we miss the human drama unfolding. There is a way of paying too much attention.

4.3.2 Ethically frivolous attention

According to the first worry, this method runs risk of running into imaginary overdrive where we start inventing things. I don't think this should worry us. No serious ethical judgement that is about widening the sphere of ethical concern seems frivolous if you think charitably about it: every extension of our imagination teaches us something new. This worry, I think, is a thinly veiled scepticism about the method in general: when you put faith in our moral imagination, you're just never going to be sure that you're getting it right. I think we've seen enough to rebut this: there is no reason to mistrust our subjective ways of engaging with the world from the outset.

The second worry, I think, is more interesting. Perhaps no widening of our ethical sphere seems frivolous in that particular case, but on the whole, we want our attention to be balanced and accorded to the right things. It seems there is something to be said for discerning ethically frivolous from ethically solid kinds of attention. There are ways of paying too much attention, and focussing on things to the exclusion of other things.

An example of this is provided by the animal-rights activists in the film *Okja* (Joon-ho, 2017). *Okja* is a film about a genetically engineered super pig that is highly critical of factory farming and ethical capitalism. Using film and narrative, following a super pig Okja and her human friend Mija, the film invites us to see that our treatment of animals in the industry is incomprehensibly horrendous. But though the film obviously sides with the animal-rights activists, it also attempts to show that animal-rights activists intense focus on animal's interests have their own problems and sometimes unpleasant sides. They tend to put animal welfare over human welfare: for example, they flirt with ideas of terrorism against human beings for their political agenda, and treat each other horribly when they suspect someone has failed the cause. They also sometimes offer up the welfare of individual animals to advance their broader goals: they knowingly send Okja to a laboratory where she is severely and painfully abused in order to capture the abuse on camera. Finally, one of the activists collapses from severe malnourishment after refusing to eat, because all food is grown with exploitation.

These activists, in their bad moments, might be seen to have lost the plot or the larger picture. This is where the worry lies, but it seems to be where the answer is too. There is such a thing as too much attention and too much sentiment in instances where it takes up such a large part of our lives that it deteriorates our attention to other creatures. But there is no reason to think thinkers and activists can't, with hard work and sincere attention, recognise these situations.

4.5 Conclusion

How do we tell the difference between a sentimental or distorting conception or way of bringing someone into focus, and one that gets it right? I outlined a racist world view espoused in Raspail's *The Camp of the Saints* (4.1). I then outlined three features of authoritative perspectives according to Cray: making sense of otherwise puzzling observations, seeing more, and emotional maturity (4.2). I argued seeing more is most useful in telling the difference between sound and mistaken judgements. One problem might be that this leads the account to have a positive bias, or that it validates ethically frivolous kinds of attention (4.3). I argue this worry isn't too worrisome.

'Seeing more' gives us a method for rejecting Raspail's racist ideology, which in turn validates the argument I made in the previous chapter: that the intuition that it is especially bad for you to be wronged or harmed in virtue of your identity.

Chapter five – Affirmative Action

So far, I have used Crary's meta-ethical approach to argue that people are moral equals, and that we should take that to be axiomatic. In this final section I will show how the view helps us to make progress on particular judgements about what particular social practices express. Taking affirmative action as an example, I will show what kind of considerations the view marks out as ethically relevant and how they can help us move the discussion forward. The main aim of the section is not to come out in favour or against affirmative action.

Affirmative action or positive discrimination describes a set of policies where oppressed groups are given preference in a selection process, for example in hiring or admission to universities. US universities have, for example, lowered admission standards for Black students, and Dutch universities have opened jobs and professorships solely to women candidates. These policies are hotly contested. Advocates argue that these policies are necessary to promote equality and redress past wrongs. Opponents argue that these policies are counterproductive and discriminatory. They are counterproductive insofar as they undermine the achievement of the targeted group, mostly benefit the most privileged people in the targeted group (high income Black people, white women), and put underqualified people in places where they can't flourish ('mismatching'). They are said to be discriminatory insofar as they distinguish between people based on legally protected characteristics like race and gender ('reverse discrimination'). Preferring Black people over white people, the thought is, is just as bad as preferring white people over Black people.

A meaning-based account of discrimination makes some progress on these questions by changing the focus. On the meaning-based account, the wrong of discrimination lies in what it expresses. Since there is no evidence to suggest these policies cast white people or men as inferior, there is no reason to think this policy is wrongfully discriminatory against them (see also Hellman, 2011, p. 79). It does, however, raise a different concern: what does affirmative action express about Black people and women? As Hellman points out, there is reason to think that affirmative action casts those it benefits as inferior: admitting Black people to universities on lower standards tracks the stereotype of Black people as intellectually inferior (Hellman, 2011, p. 79), and creating academic opportunities specifically for women tracks the idea that women aren't intellectually competitive.

I will now analyse the practice using Crary's meta-ethical method. That is, I will take seriously the subjective perspective of an affirmative action student, and explore the patterns of attention to this student that constitute moral success and moral failure. Drawing on Zadie Smith's *On Beauty* (2005) and a discussion of it by Susan Fischer (2014), I will argue that the most obvious way in which affirmative action can be demeaning is if it is done in a way that fails to pay attention to

what is important to the student, such as their interests, sense of belonging and safety. We can register this by steering away from abstract ideological ways of speaking, and towards more personal ways of speaking. The dignity of the student depends not on whether they were admitted on an affirmative action programme, but on how they are treated within this programme. I do not find evidence for the idea that affirmative action is demeaning *an sich*.

5.1 *On Beauty*: Patterns of attention

On Beauty is set in an imaginary middle-class east-coast university town of Wellington and features two politically very different families: the Belsey family and the Kipps family. The Belseys are introduced as liberal advocates of practices like affirmative action; the Kippses are Black conservative critics of it. The book follows both families as they live through a conflict on campus surrounding 'discretionary students' – Wellington university's form of affirmative action – and a simultaneous conflict in the Belsey's family life.

The first thing to note is that *On Beauty* helps us register that there are good faith ethical concerns on both sides of the debate over affirmative action. The main protagonists driving the Wellington debate on discretionary students are Howard and Zora Belsey (pro) and Monty Kipps (against). While they all view themselves as occupying the moral high ground, their real motivations are much different. All three are more interested in their own social and academic status than in the discretionary students or their work. They debate each other in a high-brow, abstract way that goes hand in hand with a failure to consider the effect their politics have on real human beings. Their insensitivity to the discretionary students is linked to two further incompetencies. First, they are emotionally incompetent. Zora does not know what she personally wants or stands for, while Howard is generally unable to relate to anyone who displays emotions. Second, they are unable to situate their personal relationships in their broader historical and cultural context (Fischer, 2014, p. 85). Howard, for example, is unable to understand the racial dimensions at work in his marriage and his betrayal of his wife. Similarly, both Howard and Monty are unable to draw on their own experiences as working-class and post-war immigration students in elite higher education. As these characters are unable to translate the historical and cultural context to the lives of real human beings, the debate turns into 'the Monty and Howard road show' (2005, p. 323).

On Beauty thus points us toward serious moral problems with *a particular way* in which affirmative action is both defended and opposed. First, these characters fail to pay attention to the wellbeing and development of the students involved. They don't consider what is important to them and

how the university could make their lives go better, nor are they able to see how they and the university are making their lives worse. This is true for both advocates and opponents. Their insensitivity does not so much tell us something about affirmative action, but rather presents us with a pattern of attention which itself exhibits a moral failure. Second, these characters fail to understand where the need for affirmative action comes from, which is the tangible racial inequality in and around Wellington. They're having the wrong *kind* of debate and are focussing on the wrong *kind* of considerations.

The novel also provides us with examples of characters who are able to see ethically relevant considerations more clearly. Poetry lecturer Claire Malcolm is able to understand racial inequality and what a place at university could mean for lower-class Black people. When her discretionary students come under fire she argues (2005, p. 160):

there are a lot of talented kids in this town who don't have the advantages of Zora Belsey – who can't afford college, who can't afford our summer school, who are looking at the army as their next best possibility [...] an army that's presently fighting a war – kids who don't –'

Similarly, Kiki Belsey is able to understand that the historical context of slavery and segregation is still affecting people's lives. She tells Monty (2005, pp. 367–368):

[...] isn't the whole point that here, in America – I mean I accept the situation is different in Europe – but here, in this country, that our opportunities have been severely retarded, backed up or however you want to put it, by a legacy of stolen rights – and to put that right, some allowances, concessions and support are what's needed? It's a matter of redressing the balance – because we all know it's been unbalanced a damn long time. In my mamma's neighbourhood, you could still see a segregated bus in 1973. And that's true. This stuff is close. It's recent.

Kiki is not an academic, and perhaps relatedly, is an emotionally mature woman who is able to understand the role race plays in her personal relationships and the lives of others. This sensitivity enables her to move the conversation away from an ideological framing, and toward the personal and historical backdrop (Fischer, 2014, p. 94). Kiki's sensitivity also enables her to do something else: she registers her personal discomfort with the position she is advocating for. In doing so she acknowledges there might be some (moral) cost to affirmative action, even if she sees the necessity and thinks the policy is good, all things considered. This attention to the ethical 'residue' is an important strength of any approach: in doing so, she is taking this dilemma *seriously*. Similarly,

Claire is an intuitive person who is ‘trying to refine and polish a ... a sensibility’, which she notes is not the same thing as having logical arguments (2005, p. 158). These, *On Beauty* argues, are the competencies and sensitivities we need to develop in order to competently speak about affirmative action.

5.1.1 Demeaning

So far, I’ve argued that there are better and worse ways to think about affirmative action. So how about our original question: is affirmative action demeaning for Black students? In a conversation with Kiki, Monty fleshes out these worries. First, Monty argues it is demeaning to tell Black young people they need special treatment, and that they are not fit for the same meritocracy as white students (2005, pp. 122, 365). He points out that Black young people are encouraged to feel like they are owed a place as a reparation for history instead of their merit. This, he argues, is a demoralising and crooked way of thinking about what they deserve and are capable of (2005, p. 365). Second, he suggests that affirmative action isn’t done with these people’s interests in mind. Instead, he argues, these Black students are being used for some larger political scheme: ‘These children are being encouraged to claim reparation *for history itself*. They are being abused as ‘political pawns’ – they are being fed lies’ (2005, p. 365).

So is Monty right? Is affirmative action demeaning? We can get clearer on this question by focusing on the character of Carl Thomas. Carl is a talented Black working-class poet and rapper who is one of the discretionary students in Claire’s poetry class. I’ve argued that it’s important to pay attention to him as a person. So what does Wellington actually do to him?

At the beginning of the novel Carl feels extremely uncomfortable in the university setting and poetry class. He acts defensively, feels like he doesn’t belong and thinks everyone is out to make fun of him and his work. This is not surprising, as Carl defines ‘large sections of his personality’ as ‘not an educated brother’ who doesn’t belong in a classroom (2005, pp. 76, 259). He has felt this way ever since his teacher had sex with him, which made all his other teachers afraid ‘he was going to rape or mug them’ (2005, p. 260). Even when he starts to engage with the class, he is ‘still not sure that this whole Wellington thing wasn’t a kind of sick joke being played on him’ (2005, p. 259). Slowly, though, Claire’s efforts get through. Carl starts to recognise that ‘Claire had that special teacher thing he hadn’t felt since he was a really small boy [...]: she wanted him to do well. [...] And he wanted to do well for her’ (2005, p. 260). Slowly, Claire is enabling Carl to view himself as ‘shoulder to shoulder with Keats and Dickinson and Eliot’ (259), and, perhaps more importantly, restoring his trust in pedagogical relationships.

As his discretionary status comes under fire, Claire pulls some strings to create a job for Carl as a Hip Hop Archivist in the Black Studies department. Carl finds pride in this work: for the first time, 'he was being hired because he knew about *this* subject, *this* thing called hip hop. [...] He had a skill, and this job required his particular skill. He was an archivist.' (2005, p. 372). At the same time, he still feels like he doesn't fully belong at Wellington. As another Black archivist tells him, it's all good and well to pretend to be a Wellingtonian, but 'people like you and me [...] we're not really a part of this community, are we? I mean, no one's gonna help us feel that way.'

Nearing the end of the novel Carl discovers Zora's real motivations for helping him (2005, pp. 414–415):

'Am I meant to be grateful? [...] apparently you wanted to do a little more than help me. Apparently you expected some payback. Apparently I had to sleep with yo' skank ass as well. That's what it was all about,' said Carl and whistled satirically, but the hurt was clear to read in his face, and this hurt grew deeper as he stumbled over further realizations, one after the other. 'Man, oh, man. Is that why you helped me? I guess I can't write at all – is that it? You were just making me look an idiot in that class. Sonnets! You been making a fool of me since the beginning. Is that it? You pick me up off the streets and when I don't do what you want, you turn on me? Damn! I thought we was friends, man!'

Carl is hurt as he comes to recognise how Zora and the other academics are using him for their own interests. He realises that 'people like me are just toys to people like you ... I'm just some experiment for you to play with' (2005, p. 418). Upset about the rejection, Zora lashes out at him: 'You think you're a *Wellingtonian* just because they let you file a few records? You don't know a thing about what it takes to belong here.' (2005, p. 417).

Carl's time at Wellington harmed him in different ways: he never fully felt part of the university, always felt like people were out to make fun of him, and was being used by people he trusted to have his best interest in mind. The pedagogical relationship he was building with Claire, his first in many years, is shattered by this realisation. The novel also points us toward other ways in which universities are hostile places for Black working-class students. Monty accuses Carl of stealing a painting from his office based solely on his background – a false accusation that could lead to serious legal problems and incarceration, and hints at what might be the alternative to university for Black American working-class men like Carl. Monty also abuses his power to have sex with a discretionary student, Chantelle Williams. He then tries to have her discretionary status cancelled to cover himself. Meanwhile, Howard is also abusing his status to have sex with a Black woman student. This sexual abuse marks universities a particularly dangerous place for Black women,

especially those whose university status is less secure and thus more dependent on their relationships with their male supervisors and teachers.

From this, I want to draw two conclusions. First, Carl and Chantelle aren't demeaned in the way Monty is worried about. The problem isn't that Carl or Chantelle didn't get in on merit, or that they are presented as someone who can't make it on their own, or that they are being cast as intellectually inferior. Even though Carl is looked down by people like Zora, as he is not an '*educated brother*' nor a '*poet poet*' (2005, pp. 76, 77), he is not a pitiful character. Carl is able to understand and analyse the world around him, and unlike Zora he knows he doesn't need the university to 'get a stage higher' with his life (2005, p. 418). He is a proud, talented person who can take care of himself, and recognises that he is emotionally, intellectually and morally more mature than many of the students and staff at Wellington. Similarly, Chantelle is not in the class 'because she is black and poor' (2005, p. 365), as Monty would have it, but because Claire thinks she is 'one of the most exciting young female talents I have come across all year' (2005, p. 160). Both Carl and Chantelle are being admitted based on talent, more so than other Wellington students from middle-class backgrounds. As Claire puts it: 'I *really* couldn't give a crap [about academic credentials]. My class rewards *talent*.'

Second, Wellington pushes Carl, Chantelle and the other discretionary students into a humiliating, undignified, insecure and unsafe position in which they get severely harmed. The pride Carl takes in his work make it even more morally jarring that he's only working there because of some political interests of other people, and that no-one at the university takes an interest in his work or what it might mean for other disadvantaged Black young people (Fischer, 2014, p. 93). As we saw before, dignity is to be found in the details of how people are treated. Moreover, Smith seems to have invented the very terminology of 'discretionary students' to underscore the marginal status these people have, and how much their status depends on those in power who can afford or deny them access (Fischer, 2014, p. 93).

As *On Beauty* shows us, the moral failure surrounding affirmative action lies in the precarity and lack of attention to what makes these students' lives go well, not in a lack of self-respect and pride, or a corrupted morale. The dignity of the student depends not on whether they were admitted on an affirmative action programme, but on how they are treated within this programme.

5.1.2 Being Used

Monty has another problem with affirmative action: he argues that Black students admitted on affirmative action programs are being used for a political agenda. Again *On Beauty* helps us see that there is a good faith ethical concern here. In *On Beauty* we see examples of people using and abusing these students for their own personal and political agenda. It is easy to see how universities can use these students for their institutional gain too: for universities, status and marketing are everything, and coming across as diverse and inclusive matters. Even though affirmative action provides opportunities some Black students otherwise wouldn't have had access to, it's not a given that the institution will care about these Black students as individuals. In many cases the opposite seems to be true: universities take in Black students and use their pictures on promotional material but fail to ensure they feel part of the university community and fail to safeguard them against racist abuse and abuse of power.²⁴ Again, the problem seems to lie in the fact that universities don't do things in the right way and for the right reasons: affirmative action is morally problematic if it's done without paying attention to what makes the lives of these students go better and worse.

5.2 Conclusion

On Beauty argues that it matters how we speak and what we pay attention to when it comes to affirmative action. To understand the necessity for affirmative action, we should look to the personal, historical and cultural context. Ideological or academic arguments that abstract from people's personal lives can obscure what is really important. Instead, moral arguments should be made by paying close attention to what makes people's lives go better or worse. From this approach, the most salient worry with affirmative action is that it isn't done with the right kind of attention for the students admitted on the programme. Done the wrong way, affirmative action puts students in a position that is undignified and unsafe. Institutions and people working in them should take special care to insulate these students from abuse of power and to provide security, both in material terms as well as in terms of a feeling of belonging. Done the right way, affirmative action enables students to take pride in their work, develop their talents and foster pedagogical relationships with their teachers and fellow students. All in all, I have found no evidence that affirmative action is necessarily demeaning. For the dignity of the students, details of their treatment matter.

²⁴ They fail, for example, to address institutional racism and sexual misconduct, and to implement functional complaints procedures (Bull, 2018; Equality and Human Rights Commission, 2019).

With this approach I've hoped to identify considerations that are ethically relevant but aren't widely recognised as such. The considerations stem from a Crary-style metaethical approach that stands out in two ways. First, the pattern of attention we evince in thinking about debates over affirmative action can already by itself constitute a moral success or moral failing. This is true for both the way in which we speak (using abstract, academic, ideological language, versus grounding arguments in personal experience, history, and culture) as well as for the way in which we pay attention to these students (do we ask them what they want and need, do we make them feel secure, are we treating them with dignity). Second, this approach gives us guidance in how to place emotionally charged and personal ethical perspectives that are likely to show up for some people in response to affirmative action. Roughly, it tells us to take those perspectives seriously: to see them as evidence that informs (and, typically, complicates) our abstract principle-driven assessments of what is ethically at stake. By contrast, an orthodox metaethical approach tells us, roughly, to be sceptical about those perspectives, since they are (presumptively) a product of subjective personal experiences, rather than objective ethical concerns.

Conclusion – No Need for Justification

Intuitively, bad things seem worse if they happen to someone in virtue of their social identity. But if that is so, what makes identity-based harms and discrimination distinctively wrong? Political and legal theorists working on discrimination and identity-based harms give us two kinds of answers. Status-based theorists, on the one hand, argue that discrimination and identity-based harms are wrong because it involves a risk to one's social status. A view like this is held by Chambers, Moreau and Hellman. Structural theorists, on the other, argue that there is no such distinctive personal wrong of discrimination. This view is held by John Gardner.

Despite their disagreements, these theorists share a common ground. They tacitly accept a picture of objectivity as excluding all reference to subjective ways of engaging with the world. On this picture, emotional and subjective responses to the world require further justification in the form of 'objective' judgements and facts. These judgements should be made 'from nowhere' and which don't depend on our sensational responses to the world; these facts should refer to a risk to tangible (status) harm. I've argued that these metaethical assumptions lead these theorists to mischaracterise the wrong of discrimination as a threat to harm, rather than a status wrong (chapter 2), and to hold the untenable positions in which they need to rely on the subjective engagement they are so suspicious of (chapter 3).

Having absorbed the Wittgensteinian metaethical ideas espoused by Crary, we might now say that the question they are trying to answer is badly-conceived. We would only go search for further justification to vindicate those moral commitments in if we had some metaethical reason to be suspicious of them. But I've argued that those intuitions don't warrant suspicion, as there's a plausible metaethical framework that can relieve us of those suspicions. On this metaethical framework subjective resources are necessary for ethics; not *prima facie* distorting. As such we can have true interpretative value judgements about when people are being treated as inferior in a way that offends against equality.

What we learn from those judgements is this: It *is* distinctively bad to be treated badly in virtue of your identity, but there's no further justificatory story for us to tell to explain or vindicate why this is the case. It is just part of the axiomatic backdrop for our moral theorising – reinforced by strong moral intuitions we have in all sorts of context – that this is the case. The intuition that bad things are worse if they happen to someone in virtue of their social identity tells us what the world is really like.

References

- Anderson, E. (1999). What Is the Point of Equality? *Ethics*, 109(2), 287–337.
- Baldwin, J. (1962). *Another Country*. Dial Press.
- Borges, J. L. (1949). Lottery of Babylon. In A. H. (translator) (Ed.), *The Aleph and Other Stories*. Penguin Classics.
- Bull, A. (2018). Silencing Students: Institutional Responses To Staff Sexual Misconduct in Uk Higher Education. Retrieved from <https://1752group.com/sexual-misconduct-research-silencing-students/>
- Chambers, C. (2004). Are breast implants better than female genital mutilation? Autonomy, gender equality and Nussbaum's political liberalism. *Critical Review of International Social and Political Philosophy*, 7(3), 1–33.
- Chambers, C. (2008). *Sex, Culture, and Justice: The Limits of Choice*. University Park: The Pennsylvania State University Press.
- Crary, A. (2016). *Inside Ethics: On the Demands of Moral Thought*. Cambridge, Massachusetts & London, England: Harvard University Press.
- Crary, A. (2018). Cognitive Disability and Moral Status. *The Oxford Handbook of Philosophy and Disability*.
- Diamond, C. (1991). The Importance of Being Human. In D. Cockburn (Ed.), *Human Beings* (pp. 35–82). Cambridge University Press.
- Diamond, C. (2016). Eating Meat and Eating People. In *Whose Peace? Local Ownership and United Nations Peacekeeping* (pp. 1–15).
- Eddo-Lodge, R. (2014). *Why I'm No Longer Talking to White People About Race*.
- Equality and Human Rights Commission. (2019). Tackling racial harassment: Universities challenged. Retrieved from <https://www.equalityhumanrights.com/en/publication-download/tackling-racial-harassment-universities-challenged>
- Fischer, S. A. (2014). “Temporal Layers”: Personal and Political History in Zadie Smith's On Beauty. In P. Tew (Ed.), *Reading Zadie Smith: The first decade and beyond* (pp. 83–96). Bloomsbury Publishing Plc.
- Fiss, O. (1976). Groups and the Equal Protection Clause. *Philosophy and Public Affairs*, 5, 107–177.
- Gaita, R. (2000). Forms of the Unthinkable. *A Common Humanity: Thinking About Love and Truth and Justice*, 1–19.
- Gardner, J. (2018). Discrimination: The Good, the Bad, and the Wrongful. *Proceedings of the Aristotelian Society*, 118(1), 55–81.
- Haslanger, S. (2012). Gender and Race: (What) Are They? (What) Do We Want Them To Be? In *Resisting Reality: Social Construction and Social Critique*.
- Hellman, D. (2011). *When is Discrimination Wrong?* Cambridge, Massachusetts, and London, England: Harvard University Press.
- Hellman, D. (2018). Discrimination and Social Meaning. *The Routledge Handbook of the Ethics of Discrimination*, 97–107.

- Joon-ho, B. (2017). *Okja*. Netflix.
- Koechlin, F. (2009). The dignity of plants, 4(1), 78–79.
- Laimann, J. (2015). Should we prohibit breast implants? Collective moral obligations in the context of harmful and discriminatory social norms. *Journal of Practical Ethics*, 3(2), 37–60.
- Loury, G. C. (2002). *The Anatomy of Racial Inequality*. Cambridge, Massachusetts, and London, England: Harvard University Press.
- McKeown, M. (2018). Iris Marion Young’s “Social Connection Model” of Responsibility: Clarifying the Meaning of Connection. *Journal of Social Philosophy*, 49(3), 484–502.
- Moreau, S. (2019). Discrimination and Subordination. In *Oxford Studies in Political Philosophy Volume 5*.
- Moreau, S. (2020). *Faces of Inequality: A Theory of Wrongful Discrimination*. Oxford University Press.
- Raspail, J. (1973). *The Camp of the Saints*. (N. (transl. . Sharpiro, Ed.). Éditions Robert Laffont.
- Sangiovanni, A. (2017). When and why is discrimination wrong? In *Humanity Without Dignity*. Harvard University Press.
- Shin, P. S. (2009). The substantive principle of equal treatment. *Legal Theory*, 15(2), 149–172.
- Smith, Z. (2005). *On Beauty*. Penguin Group. <https://doi.org/10.4324/9781315538181-14>
- Young, I. M. (2003). Political Responsibility and Structural Injustice. In *The Lindley Lecture: Freedom and Morality*. University of Kansas.