

Dehumanisation Through Slurs

Y. SANDY BERKOVSKI
Bilkent University, Ankara, Turkey

A claim sometimes made in the literature on slurs is that slurs dehumanise their targets. Here, I look at several different versions of this claim. I begin by examining and ultimately rejecting the idea that dehumanisation is a form of disrespect, and that to use a slur is to treat the target, or to think of her, with disrespect, hence, to dehumanise her. Then I look at the idea that some select slurs are aligned, in thought and in action, with dehumanising treatment in genocide. I argue that the theoretical justification of this claim is confusing, while the empirical evidence is sketchy at best. I conclude by arguing that, in general, the prospects for aligning slurs and dehumanisation are dim.

Keywords: Dehumanisation; slurs; respect; genocide; perception.

I

A claim sometimes made in the literature on slurs is that slurs dehumanise their targets. Here, I look at several versions of this claim that were defended at any significant length. In §II, I examine and ultimately reject the idea that dehumanisation is a form of disrespect, and that to use a slur is to treat the target, or to think of her, with disrespect, hence, to dehumanise her. In §III, I examine the idea that some select slurs are aligned, in thought and in action, with demonstrably dehumanising treatment in genocide. I argue that the theoretical justification of this claim is confusing, while the empirical evidence is sketchy at best. I conclude in §IV by arguing that, in general, the prospects for aligning slurs and dehumanisation are dim.

II

In this section, I address Robin Jeshion's attempt to link dehumanisation to a denial of respect. I argue that it is unsuccessful, whether or not we think of respect along broadly Kantian lines.

To explain how slurs dehumanise, Jeshion distinguishes between dehumanising thought and dehumanising action, and then also between strong and weak senses of dehumanisation. Paraphrasing only slightly, here are the four notions:¹

- (1) a. x thinks of y in a dehumanising way in the *weak* sense iff x thinks of y as a human, but an inferior one, and unworthy of full respect.
- b. x treats y in a dehumanising way in the *weak* sense iff x treats y as an inferior human unworthy of full respect.
- c. x thinks of y in a dehumanising way in the *strong* sense iff x thinks of y as someone who is not fully human (less than human), and who completely lacks the moral status of a person.
- d. x treats y in a dehumanising way in the *strong* sense iff x treats y as someone who is not fully human, and who completely lacks the moral status of a person.

Jeshion defends the weak versions (1a) and (1b). To call individuals with a slur is to think of them, and treat them, as inferior humans unworthy of full respect (Jeshion 2018: 103). Another component of her idea of dehumanisation is the characteristic emotion of contempt. Users of slurs express contempt toward their targets, and, in doing so, they regard them as “low in worth” (Jeshion 2018: 92, 94).

There are several reasons why this is a dubious view. Jeshion says that to be an “inferior human” is to have a “lower status along the dimension of *personhood* or *humanity in the moral sense*” (Jeshion 2018: 80, her italics). Now, suppose we learn that A was a ruthless murderer on a major scale who displayed few moral virtues in his public or private life. And suppose that B is a leading scientist, a committed parent, and a humanitarian. Aren’t we entitled to think of B as a morally superior human? We are indeed. We can rank people according to their moral qualities, just as we rank them according to their other qualities. If the notion of humanity is associated with certain behaviours and capacities, and if manifesting these behaviours and possessing these capacities are scalable, then being human in the moral sense is scalable too. The same goes for Jeshion’s own example (derived from Michelle Mason’s work) of a woman who comes to despise her husband (Jeshion 2018: 91). The woman now thinks less of her husband as a person. Compared to other humans, he is inferior: it turns out that he has several bad moral qualities.²

¹ See Jeshion (2018: 79). Although Jeshion repeatedly refers to David Livingstone Smith’s work on dehumanisation, his conception of it is based on essentialism, rather than denial of respect: see Smith (2011: 95ff) and also Smith (2020: 63ff). A view more in line with Jeshion’s explications is in Mikkola (2016: 164). That dehumanisation based on stereotypes (and possibly some form of essentialism) comes in degrees is argued, in a different setting, in Fiske (2021), Fiske (2011: 1–27). I return to Smith’s and Fiske’s ideas in §IV.

² To clarify, we do not here take a stance on how the putative ranking should be determined. We are rather saying that, however you think of moral value, it is

This sort of idiom is plausible, but is it at all related to dehumanisation? When we despise someone, do we thereby “dehumanise” him? I think the clear answer is no. That is not how we think or speak. By the same token, we could regard the murderer as a morally inferior human being. The woman could regard her husband as a morally inferior being. These facts aren’t sufficient for us to conclude that the individuals in question are being dehumanised. Therefore, while we may agree that to use a slur toward x is to despise x and to regard him as having low worth, these attitudes won’t amount to dehumanisation. The use of a slur may share some features with dehumanisation, but is not itself dehumanising.

One might object that, in the examples above, we have wilfully blurred the line between the justified moral inferiority of an individual and the unjustified inferiority of a whole group.³ We can rightly rank individuals based on their moral character. But by assumption, we can’t rightly rank whole groups based on some prejudice of ours, like a prejudice against gender, ethnicity, or religion. Slurs are designed to express inferiority of the latter kind. Since the two kinds of inferiority are distinct, whether we can align dehumanisation with one of them shouldn’t count as evidence for the possibility of aligning it with the other.

I reply that we are not trying to assimilate the ranking associated with, or introduced by, slurs to the moral ranking. The problem is strictly with the explication of weak dehumanisation in (1). The two kinds of ranking are indisputably different. But they are similar in that they both involve thinking of the target as inferior and unworthy of full respect. Of course, the reasons for thinking so are different in the two cases. Conceivably, that’s precisely why we may accept one kind of ranking, but not the other.

The objector might now follow two lines of argument. Both of them lead to the question of respect. On the one hand, it is tempting to tweak Jeshion’s formula (1a) just a bit:

- (2) x thinks of y in a dehumanising way in the weak sense iff:
- i. x thinks of y as a human, but an inferior one, and unworthy of full respect, and
 - ii. x doesn’t have good reasons for thinking so.

Analogously for (1b). To stick to (1a), since there are good reasons to think of A the murderer as an inferior human, thinking of him in such a way doesn’t amount to dehumanisation. But this move begs the ques-

at least plausible that you attribute it to character and actions to varying degrees. Whether you are a consequentialist or a deontologist, you are bound to think of an avaricious miser, say, as morally inferior, other things equal, compared to a generous benefactor. See, e.g., Kagan (1998: 78ff) for some discussion. It doesn’t follow though, at least according to the deontologist, that a person may have no moral status whatsoever. I return to this crucial issue shortly.

³ Slurs defined as a special kind of group pejoratives: Diaz-Legaspe (2020: 1412–1413).

tion: *Is it ever justified to think of anyone as unworthy of full respect?* For suppose it isn't. Then the condition (2ii) is vacuously true, and we are back to the same problem as earlier. The answer, therefore, must depend on the concept of respect used in Jeshion's explications.

Alternatively, one might object that, though contempt and holding in low regard are not sufficient for Jeshion's "strong dehumanisation", they are sufficient for weak dehumanisation. Weak and strong dehumanisation, says Jeshion, "align in a continuum": there is no "hard demarcation" between the two notions (Jeshion 2018: 80). At the extreme, we find rare, perfectly illegitimate dehumanisation. But you can dehumanise people to a certain degree when your act is not that rare and is perhaps legitimate. This is reflected in the situations where we say that contempt is "fitting", as may be the case with the ruthless killer and the despicable husband. And this is precisely what the concept of weak dehumanisation is intended to capture.

But is there really continuity between strong and weak dehumanisation? There is one conceptual side issue that bears mentioning. Weak dehumanisation is scalar: when you are dehumanised to this or another degree, we say that you are "weakly" dehumanised to that degree, and that you are regarded or treated as only partly human. Strong dehumanisation is Boolean: to be strongly dehumanised is to completely lack the status of a human. So to be strongly dehumanised is to be completely dehumanised. Putting the two together, we have a problem. You may be lightly or gravely injured, but you can't be "completely" injured. You may be a little dangerous or very dangerous, but not completely so. In the other direction, you may be dead, but not a little dead or very dead. We can't, that is, have it both ways: a concept can't be both scalar and Boolean. Thus, we can't speak of people dehumanised to a degree, but also of (other) people to be completely dehumanised.

It would be unwise to pin the critique on this observation alone. Jeshion might drop the talk of the "continuum" and treat weak and strong dehumanisation as two different concepts, one Boolean and the other scalar.

Hence, the main question is the same as before: what is the concept of respect used in the explications (1)? Textual evidence is inconclusive—indeed, Jeshion says nothing about her understanding of respect. One possibility is that she interprets respect in a broadly Kantian way. "Humanity in the moral sense" is meant to refer not to diverse moral qualities possessed by some humans and not by others, but rather to an intrinsic human quality *F* that separates humans from mere things. Kant's own candidate for *F* is the capacity to set goals for oneself (he also insists that this capacity, when exercised autonomously, converges on the capacity for moral action; but we needn't follow him there). Can you be an inferior or superior human in this sense? You are either able to set goals or not. Even if you don't set them very well, or even if your

goals aren't very good, you are still an end-setter. The concept of humanity so understood is not scalar, and it's not meant to be.

Further, as we see in (1a), Jeshion couples the idea of inferior humanity with "unworthiness of full respect". But if we think of moral respect along Kantian lines, then the contrast between "full" respect and "limited" or "partial" respect is impossible to draw. Under a standard interpretation, moral respect is intended to be a normative threshold we shouldn't cross in our treatment of a person, good or bad.⁴ Suppose you believe that torture is incompatible with moral respect. This means that no human, no end-setter, good or bad, should be tortured. And if you propose to torture a person just a little, or very seldom, you don't show some imperfect moral respect; rather, you fail to show moral respect altogether.

To put it another way: There are intelligible idioms of ranking people *qua* moral agents and of sorting them into "superior" and "inferior". You may have harsh words to say about a person's moral character, his standing as a moral being. You may, in so doing, also express a lot of contempt toward the person's actions and his character. But neither the idioms of ranking, nor the attitude of contempt so expressed, should have anything to do with moral respect. Vice versa, when we ask whether a certain person (or more precisely, a certain object that may in fact be a person) should be treated with moral respect, we don't mean to enquire *how much* respect is due to him, and how superior or inferior this person is.

To return quickly to (2), the rationale behind it is to look at *A*, the murderer, say, as unworthy of full respect, as a human, but an inferior one. But we now see that the attitude of respect does not scale, and that even a murderer deserves full respect. That is, the condition (2ii) is indeed vacuous, and the previous objection holds.

To prevent a misunderstanding: I'm not arguing, absurdly, that because all persons deserve respect, you can't, in practice, *de facto* deny them respect partially or wholly. I am arguing that, because of how the concept of moral respect is understood, it is conceptually impossible to deny people respect *to a certain degree*. It is impossible to say, in this broadly Kantian sense, that *A* treats *B* as worthy of some respect, but not of "full" respect. Because of the original Kantian contrast between respect and honour, respect is a matter of all or nothing.

This last claim isn't rigidly aligned with the standard interpretation of Kantian texts mentioned earlier. The same contrast between scalar and Boolean concepts is available in the accounts intended to deviate from Kant. According to Stephen Darwall, appraisal respect (meant to model "honour") may be given to a smaller or greater degree. Recognition respect (meant to model "moral respect") is "acknowledgement" of the person's dignity, the disposition to give at least some weight to the person's preferences in your practical deliberation. Self-described non-

⁴ The "standard" interpretation: e.g., Hill (1991: 11), Wood (1999: 135).

Kantians like Sarah Buss and Joseph Raz defend the notions of respect broadly in line with this formulation.⁵

It might be retorted, and I encountered this idea in discussions, that Jeshion's account is meant to describe bigoted thought, and that if such thought is shown to be incoherent, so be it. Hence, slur-users *do* dehumanise their targets, though they do that incoherently. But this response goes against some very basic demands of charity. It is not as though slur-users freely admit to dehumanising their targets. Nor does the idiom of dehumanisation belong in their vocabulary. *We* impute the concept of dehumanisation to them. By doing so, we are supposed to make certain moral qualities of their linguistic use intelligible. Therefore, if the concept of dehumanisation is shown to be incoherent, this is a failure on our part as interpreters of bigoted speech and thought. It is a sign of our failure to impute, coherently, a dehumanising thought to slur-users.

Another line of defence of Jeshion's account might run as follows. We should read the preceding argument as a *modus tollens*. Since Jeshion's account entails that respect is scalar, the notion of respect she is using is not remotely Kantian. It is some other notion of respect. Unlike Kantian respect, it is not a matter of all or nothing: it *is* scalar. Thus, consider an idea, floated by several authors, that respect is a form of care and considerateness for a particular person, sensitivity to his needs and preferences.⁶ Call this "c-respect". Let l be an interval within which it is acceptable to treat you with care. Appropriate care falls within l . There is, however, a certain threshold I shouldn't cross. If I give you too little care, let alone positively disregard your preferences, that would be unacceptable. Accorded c-respect within l is strictly normative, and a failure is a moral violation. By the same token, there are degrees of violation: we can recognise mild c-disrespect and gross c-disrespect. To put this all together, when you use a slur, you weakly dehumanise me, and when you do so, you also show me some measure of c-disrespect. We associate slur-use with weak dehumanisation, and we explicate the latter through c-disrespect.

The notion of c-respect, I agree, is a coherent one capturing some intuitions we have about "respect". But how can the use of slurs be a show of c-disrespect? How can we fail to care for people by using a slur? The only plausible answer, I think, is that slurs are offensive and insulting. Thus, according to Dillon, one way to show respect to people is by "praising and honoring" them (Dillon 1992: 109). Correspondingly, we fail to show respect when we "scorn" or "contemn" them, or when we act "arrogantly" towards them. Furthermore, it is also plausible that insult and offence come in degrees. They allow for differential moral valuation depending on their severity: slurs offend, and different slurs

⁵ Semi-Kantian respect: Darwall (2006: 122–123), Darwall (2013: 18–19, 264–270), Buss (1999: 536), Raz (2001: 159).

⁶ See, e.g., Dillon (1992), Ullmann-Margalit (2011), MacLagan (1960).

offend differently. This feeds into the proposal that, just like weak dehumanisation, c-respect is scalar.

On this reading, then, Jeshion merely wishes to point out that, since each use of a slur is mildly or grossly offensive, it is for this reason also mildly or grossly dehumanising (in the weak sense of it). There are two obvious problems with this interpretation. First, it conflicts with Jeshion's own claim, noted above, that dehumanisation is in a different league from "derogation" and "debasement" that presumably characterise insult and offence (Jeshion 2018: 80).⁷ Secondly, I take it that the rationale of linking dehumanisation to slur-use is in marking out a uniquely problematic moral aspect of slurs. Thus, ordinary offences are not supposed to dehumanise. If they were, there would be nothing special about slurs.

If dehumanisation is spelled out in terms of c-respect, however, then every offence or insult would also dehumanise, at least to some degree. Hence, slurs should lose their (im)moral uniqueness. This result would clash with the whole dialectic of Jeshion's argument. And therefore, it is not plausible to link the explications in (1) to the notion of c-respect.

I conclude that the notion of weak dehumanisation is either outright incoherent (if linked to Kantian respect) or does not serve Jeshion's purposes (if linked to c-respect). There is perhaps some other notion of respect that can be used to avoid both pitfalls. Since I can't think of any, and since Jeshion's discussion provides no clues, I won't pursue this possibility.

Can't we, finally, jettison "weak dehumanisation" altogether and stay only with the "strong dehumanisation" explicated in (1c)–(1d)? To strongly dehumanise a person in thought is to think of him as a subhuman, and to strongly dehumanise him in action is to treat him as a subhuman deserving no moral respect whatsoever. Yet this is just too strong for slurs. A speaker may use a slur and, at the same time, predicate a quality of the victim compatible with the affirmation of moral respect.⁸ The generics in (3), prejudiced and misleading as they are, are intended to attribute some bad traits to the targeted groups:

- (3) a. Wops are sleazy, but they have a great artistic taste.
 b. Chinks have no taste, but they are highly intelligent.

If these slurs dehumanise strongly, then the bigoted speakers in (3) mean to deny the slurred groups a *minimally* respectful treatment accorded to humans. That's highly implausible. It's hard to believe that those speakers are prepared to accord to typical Chinese or Italians the kind of treatment we see in concentration camps, for instance. It is similarly unlikely that the speakers in (3) even think of these groups as separate categories of non-humans or subhumans. Indeed, the mem-

⁷ The same idea underlies Jeshion's appeal to the distinction between "appraisal respect" and "recognition respect" (Jeshion 2018: 84).

⁸ The argument here echoes the one in Camp (2013) and Anderson and Lepore (2013) against the association of slurs with contempt. It is addressed in Jeshion (2018: 85–86, 96).

bers of these groups are attributed the qualities, such as taste or high intelligence, that are often thought to be uniquely human. And since the speakers profess a dislike of “wops” and “chinks” in some respects and admiration in others, it is not even clear that they mean to deny honour to members of these groups. (3) may be read as little more than expressions of a grudging admiration mixed with hostility or envy.

At this point, one might argue that the speakers in (3) could still be “quintessentialists” in the sense of Leslie (2013, 2017). They might essentialise their targets by attributing to “wops”, say, a hidden, not necessarily biological substance (“quintessence”) causally responsible for their sleaziness, i.e., “their most important, stable, and enduring” property (Leslie 2013: 110). This tendency to quintessentialise is sustained by their use of nouns and generics (Leslie 2017: 418). And it is also plausible that their quintessentialist tendency is part of the explanation why those speakers continue to entertain prejudices toward Chinese or Italians, and why they are amenable to the introduction of slurs for these groups (we don’t have to assume, though, that the extension of “wop” and “Italian” must be the same). By contrast, they don’t derogate or discriminate against accountants, to use Leslie’s own example. One reason for that is that they don’t think there is any accountant-quintessence.

Well, conceivably, the bigots in (3) do quintessentialise Italians and Chinese. But we can hardly infer from that any fact of strong dehumanisation.

For chances are that these bigots, in fact, along with us, quintessentialise many other groups identified through the same mechanism of noun use. Thus, we hear them say, “He is just a lawyer!” So, we expect them to entertain some negative attitudes, likely prejudices, toward lawyers *qua* lawyers. Unlike accountants, lawyers have quintessences, or so our bigots think. Other times, the same bigots declare, “He is a human, not a chatbot!” They mean to say that something uniquely human, the human quintessence, explains human emotional responses, say. Our bigots may well believe that. But of course, it doesn’t follow that they are in any way committed to denying lawyers or humans minimal respect. Psychological essentialism does not constitute strong dehumanisation, though it may certainly be a small step towards it, alongside many other factors.⁹

III

In this section, I examine another sophisticated attempt to tie slurs to dehumanisation that, however, is markedly different from Jeshion’s.¹⁰ At the end of §II, I appealed to the observation that Jeshion’s view fails

⁹ I explore these other factors in §IV.

¹⁰ Neufeld (2019, 2020) offers another interesting argument that I won’t be able to address here for the lack of space. Neufeld ties the use of slurs to essentialism, essentialism to the denial of free agency, which is in turn linked to dehumanisation. For an instructive early argument in the same direction see Clor (1969).

to do justice to the use of ordinary slurs like “wop” or “chink”. That was fair: according to Jeshion, all uses of slurs, certainly uses of racial slurs, are supposed to dehumanise. But what if only a few select slurs dehumanise, while others don’t?

A version of this view is defended at great length in Tirrell (2012). Taking the Rwandan genocide as her case study, she argues that certain slurs dehumanise when they are “deeply derogatory” (call such slurs “deep slurs”). Compared to ordinary slurs, deep slurs have five unique characteristics (Tirrell 2012: 190–193). Let me mention two. Deep slurs impute certain essences to the individuals designated by the term (i.e., the targets of the slurs). As observed earlier, it is questionable whether the users of “wop” or “papist” would seriously contemplate an essence common to all Italians or all Catholics. But, we now speculate, to use a slur like “*inyenzi*” (“cockroach”) or “*inzoka*” (“snake”) is precisely to attribute an essence, presumably a biological one, to the individual. Secondly, a slur like “wop”, at least in categorical subject-predicate sentences, may reliably be used to express some negative attitude toward Italians. Deep slurs, on the other hand, are not primarily vehicles for attitude expression. The presumptive deep slur, like “*inzoka*”, is “action-engendering” in the way that “wop” is not (Tirrell 2012: 176, 195). Multiple social functions are central to the use of deep slurs: “they are used to dominate, demean, or dehumanize people” (Tirrell 2012: 192). To use a deep slur is to specify which treatment is appropriate for the designated person. Moreover, sometimes the slur is used to “assign” a particular social role to the target (Tirrell 2012: 193).¹¹

Let’s try to unpack these claims. To stick to our main concern here, how exactly can we use deep slurs to dehumanise others? The answer is unusually complex and demands an extraordinary exegetical effort. Sometimes Tirrell explains the link to dehumanisation by downplaying the distinction between language and behaviour (Tirrell 2012: 206). As Tirrell says repeatedly, to use a slur is to “engender” action. It is a form of “linguistic violence” (Tirrell 2012: 176). All the same, there are two different ways for the slurs to engender action. At an early, “preparatory phase” of the genocide, the use of deep slurs “prepares” the society for physical violence (Tirrell 2012: 190, 200–201, 203, 205). More explicitly:

- (4) At an early stage of the genocide, if *S* utters “*x* is a cockroach (“*inyenzi*”)", then: (i) *S* is committed to the assertions “*x* is a pest”, “*x* is dirty”, “*x* is hard to kill”, “*x* should be killed”, and other; (ii) *S* is committed to the permission of the treatment of *x*

¹¹ Interestingly, this claim is also endorsed by Jeshion who applies it to slurs across the board, however: “By using the slur [“kike”], and by virtue of its semantics, the anti-semitic indicates that being Jewish is a *fundamental negative characteristic-defining feature* of the target’s identity *qua* person. As such, the speaker uses the slur to frame and dictate how others ought to conceive of and treat its target. If the use of the slur is public, the speaker invariably performs an act that contributes to making the property part of the target’s *social identity*.” (Jeshion 2018: 83, her italics) See also footnote 13 below.

as a pest, of the killing of x , and of certain other forms of treatment.

At this stage, linguistic use is still “dissociated from social behavior” (Tirrell 2012: 204). Let us say that at this stage, slurs have a “causal role” in dehumanisation.

The situation is different, Tirrell argues, at a later stage, when dehumanisation is accomplished by assigning particular social positions (Searlean status-functions) to the targets.

Specifically, deep slurs like “*inyenzi*” are used to assign definitive status-functions to the targets (Tirrell 2012: 196, 213). Hence, slur-use is then directly associated with harmful treatment like depriving the target of human rights. Let us say that at a later stage, slurs have a “constitutive role” in dehumanisation. Slur-use “constitutes” genocidal acts (Tirrell 2012: 185), rather than merely causing them. We may summarise this idea in the following clause:

- (5) At a later stage of the genocide: If S utters “ x is an *inyenzi*”, then S assigns to x a certain social status such that it is permissible to kill every individual having that status (same for “*inzoka*”).

We have, therefore, two different lines of argument here. Slurs dehumanise at an earlier stage of genocide by training the future perpetrators to see the victims as subhuman, and to facilitate their extermination. Slurs dehumanise at a later stage of the genocide by classifying their victims in such a way as to make them directly liable to extermination. Yet I will now argue that neither of these two arguments withstands scrutiny.

The constitutive role of slurs in dehumanisation. Clause (5) may be illustrated by the case of Simon Bikindi (214–215). At the height of the Rwandan genocide, a Hutu political agitator is visiting a prison where he identifies several Tutsis as Tutsis, who are then instantly killed. The treatment inflicted on the Tutsi victims is arguably dehumanising. I will not pause to ask why we judge it so. I will instead assume that a denial of basic legal and social rights to a person is a paradigmatic case of a dehumanising treatment.

Such a change in the legal-political status of a person may be described in a Searlean fashion as an assignment of status-functions. And these status-functions, we may agree, are dehumanising in the sense of denying minimal social roles to the victims. Yet, what does the use of slurs have to do with this? Very little. The treatment of Tutsi prisoners was dehumanising, since they lost any political status, their preferences were completely ignored, their lives counted for nothing. But addressing them with slurs did not *subject* them to that treatment. In fact, Bikindi, according to the report cited by Tirrell, didn’t even use slurs. All he had to do was to identify the prisoners as Tutsis in some way, which is indeed what he did. Of course, he could use the slurs “*inzoka*” or “*inyenzi*” instead. But their only role in the ensuing dehumanising treatment—i.e., the brutal killing devoid of any pretence of legality—was identification. The persecutors had to know who the

Tutsis were, and “*inzoka*” and “*inyenzi*” could fulfill that function well. There was nothing special about them in that regard; however, other expressions, gestures, and all sorts of indirect speech acts could do the same.

You might say: this misses the point, for how were the status-functions assigned in the first place? Didn’t slurs play a critical role in classifying Tutsis as subhuman and thus enabling their dehumanising treatment? Beyond mere identification, slurs have an action-engendering normative force: they “generate permissions” for non-linguistic behaviour (Tirrell 2012: 176). Therefore, when you say, “*A* is an *inyenzi*”, don’t you classify *A* as someone who should be killed with impunity?

I reply that the dual-content function of the slurs may be plausible. At any rate, let’s grant it: slurs serve to both identify and classify. Still, no dehumanising *treatment* should follow, nor is it even remotely likely to do so. When I say, “This is my house”, pointing at Blenheim Palace, I classify it as my house and lay my claim to it (according to Searle). But the palace wouldn’t become my house following the declaration. What is required, in addition, is making others accept my classification.¹² A host of social and political conditions have to be met for my bare classification to actually assign the intended status-function. Observe that this is a perfectly commonsensical take on what happened in Rwanda, according to Tirrell’s own narrative. It was not until 1994 that a declaration such as “*A* is an *inyenzi*” could assign the dehumanising status-function to *A*. That is because by that time, the Hutu population had *accepted* that Tutsis had to be exterminated.¹³

In short, in 1994, deep slurs could be used to identify the victims for the subsequent dehumanising treatment. This could be done by various other linguistic or non-linguistic means. Addressing *A* with a slur did not subject *A* to a dehumanising treatment, did not constitute it, and did not, *by itself*, ensure any such treatment. At times, indeed, Tirrell seems to appreciate the very special social conditions in Rwanda that could plausibly link slurs and speech in general to dehumanisation and genocide. At times, she seems to downplay them entirely. Consider this passage:

In *most* language games that use racist derogations, the derogation is a *function* attached to a person, whose status changes because of the imposition of the function. Racist derogations are status-functions, for they tell the target, “You count as a so-and-so here.” (Tirrell 2012: 213, italics added)

These claims are unconvincing. Derogation alone can’t assign any status-functions, any more than I can make a palace my house by call-

¹² Declarations and acceptance (Searle 2010: 85).

¹³ A parallel complaint should be lodged against Jeshion’s version of this idea (footnote 11). How exactly would the use of “kike” meaningfully “contribute” to the target’s “social identity”? What would such a contribution consist in, exactly? And wouldn’t this require a class of non-trivial social conditions to obtain, rather than just the bare use of a slur?

ing it “my house”. For the function to be assigned, it must gain widespread acceptance and become an institutional fact. In most cases of derogation, there isn’t any such acceptance.

I conclude that there is no good argument to show the constitutive role of deep slurs. However, a large question still remains, and it is whether their use played a unique causal role in creating the conditions for a future dehumanising treatment. Through their systematic use, for example, prejudices may gradually become entrenched, thus paving the way for dehumanising treatment. This is the question of their role in dehumanising thought at the early stages of the genocide, and I turn to it next.

The causal role of slurs in dehumanisation. The predication “A is a snake” may be understood literally or metaphorically. If understood literally, then *A* is said to be a member of the animal kingdom having a particular appearance, mobility, and other physical characteristics, including the manner in which it can be killed most efficiently. If understood metaphorically, *A* may well be human but possess some snake-like characteristics: attacking quickly, escaping dangerous situations, moving stealthily, and poisoning reputations. Only some of these characteristics may be predicated of animal snakes. Others will be only metaphorically similar (“poisoning reputations”) to those of animal snakes. At any rate, using this predicate does not license an inference to the conclusion that the metaphorical snake should be literally killed, let alone killed in some specific snake-like way.

If we accept Tirrell’s idea in (4) about the link between assertion and action, then the difference between the literal and metaphorical uses may be stated as:

- (6) a. If a speaker *S* asserts on the occasion *O* that *x* is metaphorically *F*, then: (i) *S* is committed on *O* to ascribe to *x* some properties *G1*, *G2*,... that are widely ascribed to *F*s in *S*’s linguistic community and are salient on *O*; (ii) *S* is committed to treat *x* the way *G1*s, *G2*s, ...are treated in *S*’s community.
- b. If a speaker *S* asserts on the occasion *O* that *x* is literally *F*, then: *S* is committed on *O* to treat *x* the way *F*s are treated in *S*’s community.

Now, in the course of the Rwandan genocide, was the use of “*inyenzi*” and “*inzoka*” metaphorical or literal? On the one hand, Tirrell explicitly says that the use of “*inyenzi*” was in fact metaphorical (Tirrell 2012: 200–201; also compare the idiom of “depiction” on page 203). But then the proposed explanation does not add up. To repeat: When you say, “Harry is a pig”, your utterance, insulting as it may be, does not commit you to feeding Harry the leftovers or delivering him to the slaughterhouse. Certainly not: Harry is only metaphorically a pig. When, in the Rwandan context, a speaker says, “*A* is an *inyenzi*”, *A* is, again, only metaphorically a pest. Yet, according to (4), the utterance inexplicably commits you to killing *A* like any other (literal) pest. This is *ad hoc*. *A*

metaphorical use of the predicate “*x* is an *inyenzi*” should be like any other metaphor. Then its use is no more dehumanising than the use of “*x* is the sun”, “*x* is a rat”, “*x* is a lion”.

Other revealing details in Tirrell’s account suggest, however, that at least for some speakers, the use of deep slurs was *not* metaphorical.¹⁴ Consider this crucial passage about how the inference was actually made:

The extremist propaganda *described* Tutsi as cockroaches or snakes. For many *uneducated peasants*, if the official authorities state that Tutsi are snakes, *it can’t be wrong*. If the local official of the commune orders people to kill snakes, *it makes sense*. When you kill a snake, you smash its head, then you cut it up in different places to make sure it’s really dead. These very same forms of torture were inflicted on many Tutsi. (Tirrell 2012: 205, citing a report by André Sibomana, italics added)

A plausible reading of this report is that the Hutu peasants understood the predicate “*x* is an *inzoka*” literally. They took it to have plain descriptive content. Because Tutsis were literal, animal snakes, they had to die a snake-death, rather than a more familiar human death. This literal understanding was not doubted despite its evident absurdity, since it came from the authority.

But then the explanatory challenge should lie elsewhere. If I come to believe that you are a fly, don’t ask why I spray you with a repellent: that’s what people do to flies (rightly or wrongly). The challenge would rather be to explain how I came to have this preposterous belief about you in the first place. In the Rwandan situation, at least if we take Sibomana’s report at face value, the challenge is to explain how the peasants could trust the authority so blindly as to develop such a radical delusion. Potentially, the use of slurs could be a causal factor in the treatment of the Tutsi victims if one could demonstrate that it was uniquely suited to trigger the delusion. Of course, Tirrell’s argument invokes no such evidence. On the face of it, the Hutu peasants merely understood “*x* is an *inzoka*” as a descriptive term.

Supposing, however, that slurs used by the authority *could* cause that radical delusion, should we *then* describe the resulting behaviour as “dehumanising”? This kind of scenario provokes conflicting intuitions. Compare insults: you mistake your elderly boss for a friend and approach him with “Hey, dude!” The boss may *feel* insulted. But once the facts are known, not only did you not intend to insult him, but there was arguably no insult to begin with.¹⁵

As I said, intuitions vary. At a minimum, in the Rwanda situation, we should only insist on the idea that dehumanisation, if there is any, does not explain why the killings happened. Furthermore, there is at least one theory of dehumanisation that delivers an unambiguous result. According to Smith (2020, 2021), a radical delusion of the kind

¹⁴ See Smith (2021: 18ff) for an extended statement. I address Smith’s view in §IV.

¹⁵ For a useful discussion on insults see Milić (2018).

apparently experienced by the Hutu peasants is, on its own, irrelevant to the fact of dehumanisation. I return to this issue in §IV.

The upshot is a dilemma. If the use of deep slurs was metaphorical, then no specifically non-human properties were attributed to the target. No properties were attributed, at any rate, that could, *in themselves*, amount to thinking about the target in a dehumanising way. Moreover, if there were such properties, then the use of ordinary metaphors would be equally dehumanising. This is a cost we should be unwilling to pay. If, by contrast, the use of those predicates was literal, then there was a tragic delusion. Hutu peasants killed their victims, since they came to think of them as real, animal vermin and real snakes. Here we can explain the treatment of the victims, but so far as it was a result of a delusion, dehumanisation, if there was any, remains explanatorily idle.

Now, if we embrace the first horn of the dilemma, we must say that the Hutu speakers did not at all think of the targets as subhuman or non-human. If we embrace the second horn of the dilemma, we must say that the speakers did think of them as non-human, but only because of a delusion prompted by their trust in the authority. This explanation, although possible, does not generalise and is plainly irrelevant to our purposes. Ultimately, therefore, Tirrell's account doesn't help to align deep slurs either with dehumanising thought (at an early stage of the genocide) or with dehumanising treatment (at a later stage of the genocide).

IV

So far, my discussion has been negative. I have argued that the views examined in §§II–III fail to offer good reasons for either a constitutive or a causal link between the use of slurs and dehumanisation. It would be fitting to conclude with an argument that shows, in principle, why we can't find such links. One way to proceed is to offer a theory of slurs' meaning and then explain why slurs have little to do with dehumanisation. This I can't do: given the very large literature on both subjects, it would at least require another long paper. For the same reason, I won't be able to offer a theory of dehumanisation. Instead, I will claim that, within the debate on slurs, the notion of dehumanisation was invoked for certain rhetorical purposes. But the notion that could serve those purposes cannot plausibly be aligned with the use of slurs, whatever our position on their semantics and pragmatics.

Why would slur theorists insist that the use of slurs is dehumanising? What was their reason for using the idiom of "dehumanisation" specifically? Call this broad claim the "dehumanisation thesis". Conceivably, their motivation was to charge slur-users with a grave moral violation. This goes pretty much without saying, and explicit statements are hard to find. Still, recall Jeshion's account in §II. Jeshion asks whether terms like "derogate" or "degrade" could adequately de-

scribe the linguistic function of slurs. In her view, a slur-user merely regards the target of slurs as an inferior human, rather than downright subhuman (the case of strong dehumanisation). But precisely because derogation and dehumanisation are “continuous” with each other, we must insist on the idiom of “weak dehumanisation” (Jeshion 2018: 80). In general, the writers we have discussed wouldn’t have invested a major effort in showing that slurs dehumanise, had they not assumed, and wished their readers to assume, that dehumanisation is a very serious moral wrong. Call this the “moralistic assumption”.

Hence, the concept that slur theorists are interested in (call it “m-dehumanisation”) must satisfy this condition:

(Gravity) M-dehumanisation is a moral violation of the highest degree.

Its next feature is a corollary of the first one (already noted in §II):

(Rarity) M-dehumanisation is relatively rare.

A grave moral transgression can’t happen too often. How often is too often? Plainly, no general formula is possible. Still, a grave transgression can’t be routine. For example, some authors believe that many phenomena of modern culture are dehumanising. In Montagu and Matson, dehumanisation is interpreted as “the victory of the mechanical equation over the organic principle” (Montagu and Matson 1983: xxix). As a result, people are “emptied out of essential humanity—in order to be restocked with artificial needs and scientifically conditioned reflexes.” Punk rock, Martin Scorsese’s movies, professional sports, Marcel Duchamp’s art—all of them dehumanise by representing people as, or encouraging them to be, mechanical and animalistic. But even though these authors diagnose certain social and cultural activities as a malaise, they don’t intend to stop or ban them. Clearly, they don’t condemn artists or athletes as perpetrators of a terrible moral transgression.

Thus, if nearly all areas and cultural processes of modern life are dehumanising, as per Montagu and Matson, then this fact alone should diminish the scale of their putative immorality. Certainly, there will be conditions, as in totalitarian and genocidal regimes, that would enable a routine perpetration of moral crimes. But then, in the larger historical context, these conditions should themselves be relatively rare. If, for example, you are told that all of humanity is stained by an original sin, then the immorality of that bare sin can’t be too high.

Similar considerations favour another feature:

(Complexity) M-dehumanisation is a relatively complex task to accomplish.

Again, we can’t tell in advance how complex it is. But we may agree that something as simple as moving your finger can’t m-dehumanise. Easily performed actions would violate (Rarity).

If such are the constraints on the concept of m-dehumanisation, what are the implications? Well, in the first place, there can’t be m-dehumanising thoughts. They can’t plausibly meet any of the three constraints. No bare thought, however morally repugnant, should count as

a *grave* moral transgression. The same applies to rarity and complexity. Thus, from the outset, only a specific treatment could be considered m-dehumanising. And if the use of slurs does not constitute such a treatment, then it can't properly be described as "m-dehumanising".

This conclusion shouldn't be too controversial. Paradigmatic instances of dehumanisation that satisfy our three criteria would be certain kinds of treatment inflicted on humans, like imprisonment in labour camps, genocidal violence, and possibly some instances of extreme poverty.

Conceivably, that is why dehumanisation is associated with a grave moral wrong to begin with. If the present view is at all controversial, it is only because we now say that these forms of treatment are, in fact, the *only* proper instances of dehumanisation (interpreted as m-dehumanisation).

Perhaps, however, our imputation of the moralistic assumption is uncharitable, and the dehumanisation thesis has some other reasons behind it. Perhaps the authors we have explored here claim that slurs dehumanise only to highlight specific qualities of thought characteristic of slur-users. Further, they are entitled to call those qualities "dehumanising" not because of their immorality, but rather because the same qualities are recognised as dehumanising in other areas of research, in particular in social psychology, where the focus is on "dehumanising perception". As before (§III), the claim may be either about all slurs or just some of them:

- (7) a. For every slur *s*, for every speaker *x*, if *x* competently uses *s*, then *x* has a dehumanising perception of the target of *s*.
- b. There is a slur *s* such that, for every speaker *x*, if *x* competently uses *s*, then *x* has a dehumanising perception of the target of *s*.

There are now two approaches in the psychological literature. According to one approach, primarily associated with David Livingstone Smith, dehumanising perception is Boolean. It has the following condition:

- (8) *A* has a dehumanising attitude toward *B* iff *A* sees *B* as a sub-human animal. (Smith 2020: 19)

As Smith argues at length, dehumanisation is a unique phenomenon (Smith 2021: 26ff). It is not continuous with other derogatory attitudes. In particular, to think of *x* as subhuman is not to attribute to *x* a certain descriptive property, like the lack of rationality or of emotions. It is rather to stipulate that *x* is not "one of us": having first stipulated that he belongs to a natural kind of "humans", the speaker next stipulates that *x* does not belong to that kind (Smith 2021: 155). This entails, Smith further argues, that dehumanising thought is politically significant (Smith 2021: 157). It is part of the "ideology" designed to "produce, perpetuate, and augment oppression" (Smith 2021: 158).

One immediate benefit of Smith's view is that it resolves the tension surrounding radical delusion (§III). As a result of their (putative) delusion, the Hutu peasants failed to see their Tutsi victims as human

at all. And they thought so precisely because they attributed to them a certain descriptive property of being non-human, on a par with the property that snakes or insects might have. In contrast to the properly dehumanising thought, they didn't form that belief on the grounds of first thinking of Tutsis as "not the same as us". Further, their delusion itself was not part of what Smith calls "ideology"—though it was very likely a consequence of that oppressive ideology promoted by the authorities they trusted.

It seems now that only (7b) is a live option: only some slurs, like animalistic slurs "*inzoka*" or "*Judensau*" (discussed by Smith), may be associated with dehumanisation. Yet even (7b) may be too strong. It is a claim about the semantics of animalistic slurs. On many accounts, it should be possible to use a slur competently without entertaining anything like a dehumanising thought in (8).

In fact, Smith himself is apt to reject (7b). The dehumanising thought associated with the slur *s* is part of ideology. But Smith argues that ideologies can remain latent for a long time. This can mean, for example, that at some point in medieval times, the slur "*Judensau*" could be used as an instrument of oppression (to mark out those who are not part of society), given the political conditions at the time and the willingness of the population to persecute the Jews. Later on, as conditions change, the same slur may be used when it is not, in fact, part of the oppressive ideology.

As the conditions change further, the slur can once again become part of the ideology. Then, only for those people who endorse this ideology and participate in its oppressive practices, the use of that slur should be associated with dehumanisation. To give a not-too-unrealistic example, imagine a German émigré in California in 1939 who disapproves of Nazi racial persecution. The émigré has his prejudices: in fact, he is antisemitic. Yet he thinks that no dislike of the Jews, however plausible, should be turned into a state policy. On the lips of this person, the use of "*Judensau*" will not be dehumanising. At the very same time, on the lips of a Nazi official fully on board with racial persecution, its use *will* be dehumanising.

If this reasoning is sound by Smith's lights, as I think it is, then instead of (7b) we should only accept its watered-down version:

- (9) There is a slur *s*, there is an occasion *o*, there is a speaker *x*, such that if *x* competently uses *s* on *o*, then *x* has a dehumanising perception of the target of *s*.

The occasion *o* is special: many political factors must be in place to enable exclusion and persecution of *s*'s target. Since the bare use of a slur does not create an ideology and persecution, it does not warrant an ascription of a dehumanising perception. Hence, Smith's account can't be a good reason for the dehumanisation thesis: it warrants a link between slurs and dehumanising perception only on some rare occasions, and only for some slurs. Thus, in the case examined by Tirrell, we can certainly attribute dehumanising perceptions to Simon Bikindi and his

ilk. But using “*inzoka*” and “*inyenzi*” is just one way to express these perceptions, and it is not a necessary condition for these perceptions to arise. Equally, as already argued earlier, dehumanising treatment does not require the use of any dedicated slurring terms.

As far as the theory of slurs is concerned, Smith’s contention comes down to this: combined with many other political factors, dehumanising thought interpreted along the lines of (8) enables atrocities and violent persecution. Some of these actions and policies we describe as “dehumanising”. Slurs, however, have a marginal role to play there. Only under very specific, historically unique conditions can we attribute dehumanising thought to its users. In contrast to Tirrell’s view, the use of slurs can’t constitute a dehumanising treatment, either. I conclude that Smith’s account doesn’t give us a good reason for any interesting link between slurs and dehumanisation.

Smith’s view is an outlier. Most researchers in this debate favour a scalar notion of dehumanisation. There are several versions of this idea. Haslam (2006, 2020) suggested a “dual model” of dehumanising attitudes. Humanness should be interpreted along two dimensions. Cognitive abilities and self-control distinguish humans from animals. Sympathy, warmth, and emotions in general distinguish us from machines and robots. Thus, dehumanisation in its extreme forms is “animalistic”. It is characterised by a vertical comparison: humans are higher, better creatures than the animal-like subhumans. It is also aligned naturally with disgust and contempt (Haslam 2006: 258). “Mechanistic” dehumanisation, in which certain groups are portrayed as machines, is milder. It is characterised by a horizontal comparison: machine-like individuals are seen as strange, emotionally deficient, but not necessarily inferior. Hence, you might appreciate the ruthless efficiency of a “mechanical” person. The effect typically triggered by this form of dehumanisation is indifference, rather than disgust.

A similar idea is in the “stereotype content model” of cognition in Fiske (2011, 2021).¹⁶ Let’s examine it in more detail. In their interactions, we are told, people divide social groups along the dimensions of competence, on the one hand, and trustworthiness and friendliness (or “warmth”), on the other. One’s own in-group is usually praised both as competent and trustworthy. Elderly people are pitied as incompetent, but trustworthy. Successful out-group members (e.g., Asians in the US) are envied as competent but untrustworthy. Finally, homeless people, members of the “extreme out-group”, are perceived as disgusting and are scorned as both incompetent and untrustworthy.

¹⁶ Another closely related idea is developed in the theories of “infracumanisation”, where A is said to infracumanise B iff A thinks of B as lacking secondary emotions like remorse, pride, or nostalgia, that, according to A, characterise humans. See Leyens et al. (2007) and Rodríguez-Pérez and Betancor (2023) for a review. Although the terminology of “infracumanisation” is overtly chosen to contrast it with “dehumanisation”, the comments I make below about Fiske’s view apply to infracumanisation theories as well.

This model fits very naturally with the idea that slurs are linked to dehumanisation. We can, for example, paraphrase Jeshion's conditions (1) and say:

- (10) x targets y with a slur s iff x thinks of y as belonging to an out-group G such that members of G are lacking either in competence or in trustworthiness, or both.

Each slur will have an identification function: it will assign y to a particular out-group G . The right-hand side of the biconditional is what, on this view, constitutes dehumanisation. And, to repeat, when we say that it does that, we are not spelling out a philosophical analysis of dehumanisation. Rather, we simply borrow idioms from the psychological research, like Fiske's.

The problem with this manoeuvre is that Fiske's model seems motivated by the same moralistic assumption we have criticised earlier. That people divide social groups in accordance with certain (alleged) properties is an intriguing generalisation ostensibly supported by empirical evidence. But, Fiske insists further, the categorisations of out-groups are also "dehumanising". Individuals in the extreme out-group are "extremely" dehumanised, the rest are dehumanised to a lesser extent. This is because to dehumanise is to "refuse to acknowledge the complex human experience of other people" (Fiske 2011: 21). To different degrees, we deny them warmth, familiarity, emotions, articulate language, and complex minds.

Note the implications. If you faithfully and patiently care for your elderly grandfather, whom you categorise as "elderly" and who, you believe, no longer has the best command of his mental faculties, then your attitudes are now said to be "dehumanising", at least to a degree. Is there anything wrong with this description? Or rather: why should we choose exactly the idiom of dehumanisation to describe those kinds of attitudes?

I suspect that the choice, again, has to do with rhetoric. We begin with instances of dehumanising extreme out-groups that, on the face of it, are plausibly associated with a moral violation. Then we insist on the similarities between the attitudes toward these groups and the much broader attitudes toward many other groups. Many more attitudes, and many more people, are now implicated in essentially the same kind of violation. The risk here is not only the questionable expansion of moral violations. It is also that the paradigmatic, really serious violation is no longer grave or unique. For nearly everybody is doing the *same thing*, we say, just to another degree.¹⁷

What of the scorn and disgust toward extreme out-groups, the supposed primary locus of dehumanisation? We are told that:

Dehumanized perception is a cognitive bias characterized by *spontaneous* failure to think about mental contents—thoughts and feelings in a social target's mind. [...] People *spontaneously* fail to consider another person's

¹⁷ For a similar complaint see Bloom (2022).

mind if that person elicits predominantly disgust, a phenomenon we characterize as dehumanized perception. [P]articipants *spontaneously* think about the minds of dehumanized targets less often than the minds of other social targets. (Harris and Fiske 2011: 175, 180, italics added)

Suppose that this cognitive bias is real. Should we call it “dehumanising”? If this is just another label, it’s as good as any. Yet it can’t be just another label: its use is intended to strike a chord. We are expected to agree that there is something deeply repugnant about that cognitive failure. This is far from clear. The problem is not only that we should hesitate to judge a “spontaneous” failure, over which a person has no immediate control, as immoral. It is also that merely having certain experiences, let alone spontaneous ones, is not a moral failure. Or if it is, we can’t grant it without further argument.¹⁸ It would be different if, in addition, we were shown how these experiences lead to an “m-dehumanising” treatment—that is, to a treatment that can straightforwardly be interpreted as morally repugnant.

I think no evidence for this causal link is forthcoming. For we are not looking for *some* causal link between spontaneous biases of the kind described by Fiske and m-dehumanising treatments like various kinds of persecution. The truth is that many causal factors enable these treatments. Consider, for example, the Holocaust. In the first place, there were transparent political and economic factors that provided ample motivation for murderous actions.¹⁹ Secondly, as many historians have argued, the Holocaust was enabled in large part by the efficient German bureaucracy, as well as the German military and its highly distinctive ethos.²⁰ These social actors played an incomparably more central role in the Holocaust than Alfred Rosenberg’s racist ramblings. It would, however, be strange to call those institutions and their practices “m-dehumanising” merely because of their causal role (though they may perhaps be mechanistically dehumanising in Haslam’s sense of the term). So the evidence we should look for must show how biased thinking leads directly and inescapably, pretty much on its own, to m-dehumanising treatment. But since this treatment results from a combination of complex social factors, we are unlikely to find such evidence.

Summing up, there is no good reason to characterise the use of slurs as a form of dehumanisation. A likely explanation for why different writers continue to insist on linking slurs to dehumanisation is their attempt to charge the users of slurs with a grave moral violation (the

¹⁸ For an extended argument why thinking can’t be morally wrong, see Sher (2021).

¹⁹ Typical economic factors: Evans (2004: 22ff). Conspiracy theories: Evans (2009: 242ff).

²⁰ German bureaucracy and the Holocaust: Hilberg (2003: 1059ff). German military and the Holocaust: Elias (1996). The complexity of the more remote causes of the Holocaust, like the tendency to blind obedience, is addressed in Evans (1987: 156ff).

“moralistic assumption”). But whatever moral violation can plausibly be aligned with the use of slurs, it is simply not grave enough to count as dehumanisation. The less likely explanation is that those writers merely aim to join ranks with social psychologists who have developed theories of “dehumanising perception”. But in this case, either those theories, like Smith’s, are not altogether relevant to the use of slurs, or else they rest on the same problematic moralistic assumption.²¹

References

- Anderson, L. and E. Lepore. 2013. “Slurring words.” *Noûs* 47 (1): 25–48.
- Bloom, P. 2022. “If everything is dehumanization, then nothing is.” *Trends in Cognitive Sciences* 26 (7): 539.
- Buss, S. 1999. “Respect for persons.” *Canadian Journal of Philosophy* 29 (4): 517–550.
- Camp, E. 2013. “Slurring perspectives.” *Analytic Philosophy* 54 (3): 330–349.
- Clor, H. M. 1969. *Obscenity and Public Morality*. Chicago: University of Chicago Press.
- Darwall, S. L. 2006. *The Second-Person Standpoint*. Cambridge: Harvard University Press.
- Darwall, S. L. 2013. *Honor, History, and Relationship*. New York: Oxford University Press.
- Diaz-Legaspe, J. 2020. “What is a slur?” *Philosophical Studies* 177 (5): 1399–1422.
- Dillon, R. 1992. “Respect and care.” *Canadian Journal of Philosophy* 22 (1): 105–131.
- Elias, N. 1996. *The Germans*. New York: Columbia University Press. Edited by M. Schröter. Translated by E. Dunning and S. Mennell.
- Evans, R. J. 1987. *Rethinking German History*. London: Allen and Unwin.
- Evans, R. J. 2004. *The Coming of the Third Reich*. New York: Penguin Press.
- Evans, R. J. 2009. *The Third Reich at War*. London: Penguin Books.
- Fiske, S. T. 2011. *Envy Up, Scorn Down*. New York: Russell Sage Foundation.
- Fiske, S. T. 2021. “How status and interdependence explain different forms of dehumanization.” In M. E. Kronfeldner (ed.). *The Routledge handbook of dehumanization*. London: Routledge.
- Harris, L. T. and S. T. Fiske. 2011. “Dehumanized perception.” *Journal of Psychology* 219 (3): 175–181.
- Haslam, N. 2006. “Dehumanization: An integrative review.” *Personality and Social Psychology Review* 10 (3): 252–264.
- Haslam, N. 2020. “The many roles of dehumanization in genocide.” In L. S. Newman (ed.). *Confronting Humanity at Its Worst*. New York: Oxford University Press, 119–138.
- Hilberg, R. 2003. *The Destruction of the European Jews*. New Haven: Yale University Press.

²¹ I am grateful to two anonymous referees for helpful comments on an earlier draft of this paper.

- Hill, Jr., T. E. 1991. "Servility and self-respect." In *Autonomy and Self-respect*. Cambridge: Cambridge University Press, 4–18.
- Jeshion, R. B. 2018. "Slurs, dehumanization, and the expression of contempt." In D. Sosa (ed.), *Bad Words*. New York: Oxford University Press, 77–107.
- Kagan, S. 1998. *Normative Ethics*. Boulder: Westview Press.
- Leslie, S.-J. 2013. "Essence and natural kinds." In T. Szabó Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*. Oxford: Oxford University Press, 108–165.
- Leslie, S.-J. 2017. "The original sin of cognition." *Journal of Philosophy* 114 (8): 393–421.
- Leyens, J.-P., S. Demoulin, J. Vaes, R. Gaunt, and M. P. Paladino. 2007. "Infra-humanization: The wall of group differences." *Social Issues and Policy Review* 1 (1): 139–172.
- MacLagan, W. G. 1960. "Respect for persons as a moral principle: I." *Philosophy* 35 (134): 193–217.
- Mikkola, M. 2016. *The Wrong of Injustice*. New York: Oxford University Press.
- Milić, I. 2018. "What counts as an insult?" *Acta Analytica* 33 (4): 539–552.
- Montagu, A. and F. Matson. 1983. *The Dehumanization of Man*. New York: McGraw-Hill.
- Neufeld, E. 2019. "An essentialist theory of the meaning of slurs." *Philosophers' Imprint* 19 (35).
- Neufeld, E. 2020. "Pornography and dehumanization." *Australasian Journal of Philosophy* 98 (4): 703–717.
- Raz, J. 2001. *Value, Respect, and Attachment*. Cambridge: Cambridge University Press.
- Rodríguez-Pérez, A. and V. Betancor. 2023. "Infrahumanization." *Current Opinion in Behavioral Sciences* 50: 101258.
- Searle, J. R. 2010. *Making the Social World*. New York: Oxford University Press.
- Sher, G. 2021. *A Wild West of the Mind*. New York: Oxford University Press.
- Smith, D. L. 2011. *Less than Human*. New York: St. Martin's Press.
- Smith, D. L. 2020. *On Inhumanity*. New York: Oxford University Press.
- Smith, D. L. 2021. *Making Monsters*. Cambridge: Harvard University Press.
- Tirrell, L. 2012. "Genocidal language games." In I. Maitra and M. K. McGowan (eds.), *Speech and Harm*. Oxford: Oxford University Press, 174–221.
- Ullmann-Margalit, E. 2011. "Considerateness." *Iyyun* 60: 205–244.
- Wood, A. W. 1999. *Kant's Ethical Thought*. Cambridge: Cambridge University Press.