



This is a repository copy of *Are Generics Especially Pernicious?*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/110758/>

Version: Accepted Version

Article:

Saul, J.M. (2023) *Are Generics Especially Pernicious?* *Inquiry*, 66 (9). pp. 1689-1706.
ISSN 0020-174X

<https://doi.org/10.1080/0020174X.2017.1285995>

This is an Accepted Manuscript of an article published by Taylor & Francis in *Inquiry* on 13 Feb 2017, available online: <http://www.tandfonline.com/10.1080/0020174X.2017.1285995>.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Are Generics Especially Pernicious?¹

In recent years, both within and outside philosophy, the idea has taken hold that generics play a special role in perpetuating social injustice. Sarah-Jane Leslie's work (Leslie forthcoming; Rhodes, Leslie and Tworek 2012; Wodak, Leslie and Rhodes 2015), and Sally Haslanger's building on it (Haslanger 2011), had led to an increasingly widespread belief that the use of generic expressions plays a crucial role in passing on and shoring up racist and sexist beliefs. I have started to hear feminists catching themselves using generic terms to describe social groups, and trying to rephrase their utterances at conferences, citing Haslanger's and Leslie's work as the reason. But, as I will explain here, I think we do not yet have good reason to think that we should single out generics about social groups out as peculiarly destructive, or that we should strive to eradicate them from our usage. Indeed, I think they continue to serve a very valuable purpose and we should not rush to condemn them.

I view the current attention to generics as a mistake in the battle against prejudice. This is not because they pose no problems—they do pose problems (though I am not yet convinced about many of the problems claimed by Leslie and Haslanger). But so do *many* other closely related constructions. If we focus our energies on avoiding generics, substituting other phrases that we (wrongly) take to be innocuous, we will be making a serious error. Instead, we need to think much harder about how to confront many sorts of utterances which make reference to social groups.

1. Striking Property Generics

The focus of my discussion will be mainly on the generics that Leslie dubs “Striking Property Generics”. However, some of the concerns raised will be relevant to other generics as well. Striking property generics are especially puzzling, and (we will see) especially politically interesting. Consider (1) and (2) below, from Leslie (forthcoming):

- (1) Mosquitoes carry West Nile Virus.
- (2) Books are paperbacks.

Intuitively, (1) is true and (2) is false. But on reflection these intuitions are very puzzling: only a

¹ This paper was for some time delivered and circulated as “Generics don't Essentialise People; People Essentialise People!”.

tiny minority of mosquitoes (less than 1%) carry West Nile Virus, and the overwhelming majority of books (around 70%) are paperbacks. Leslie calls generic claims like (1)—which attribute particularly dangerous properties — ‘Striking Property Generics’. Striking property generic claims are readily accepted on the basis of just a few instances, and often just one.

Importantly, however, not all claims which attribute dangerous properties are so readily accepted. Learning that an office chair exploded with fatal consequences (<http://unusualdeaths.com/2014/03/23/boy-killed-when-office-chair-exploded/>) does not make people inclined to accept (3):

(3) Office chairs explode.

Yet learning about a single dramatic act of terrorism by a Muslim might well (as Leslie notes) lead people to accept (4):

(4) Muslims are terrorists.

So how does this work?

According to Leslie, one necessary condition, both for people’s acceptance of a Striking Property Generic claim and for its truth, is that being a member of the kind must be a “good predictor” of having the dangerous property attributed. Here is how she cashes that out:

“It matters, then, for the truth of ‘mosquitoes carry the West Nile Virus’ that the virus- free mosquitoes will carry the virus if circumstances allow. ‘Sharks attack bathers’ is true only if the sharks that never in fact cause harm to humans would typically do so given half a chance, and so on.” (Leslie forthcoming: 15)

Nobody would ever think that office chairs have a disposition to explode “if given half a chance”, so this nicely explains why we’re not tempted by the claim “office chairs are dangerous”.

But, Leslie notes, we don’t have very good direct access to what many kinds of things are disposed to do. So, she suggests, we rely on various heuristics. One thing we do is to look for kinds for which we think that “there is some hidden, non- obvious, and persistent property or underlying nature shared by members of that kind, which causally grounds their common properties and dispositions.” (Leslie forthcoming: 17) We do think that mosquitoes are like this. We don’t think that *flying things* are, which explains why we’re much more likely to accept that mosquitoes carry West Nile Virus than that flying things do. Those kinds that we think of as

having such a shared nature are what Leslie calls “essentialised kinds”, and social groups that we take to have an essence are essentialised social groups.

We only accept striking property generics as true, Leslie suggests, when the kinds or social groups they are about are essentialised. However, this is not sufficient for their truth. For striking property generic claims to be true—rather than just commonly accepted as true—*it actually needs to be the case* that being a member of the kind is a good predictor of possessing the property. And this, in turn, most likely requires that the kind genuinely be essentialised. Now we are in a position to understand Leslie’s story about why both (3) and (4) are false, and why people make errors in thinking about (4).

(3) Office chairs explode.

(4) Muslims are terrorists.

Both of these claims are false because being a member of the kind is a poor predictor of the property in question. According to Leslie, a key to people’s *belief* that (4) is true is that they take Muslims to be an essentialised kind. They use membership in an essentialised kind as a proxy for deciding whether kind membership is a good predictor of the property, so these people wrongly take (4) to be true. In the case of (3) we are not tempted to make such errors. First, we don’t think of office chairs as an essentialised kind—while they share many properties, these properties are all rather obvious, rather than hidden. Next, we have so much experience with non-exploding office chairs that we are not even for a second tempted by the thought that being an office chair is a good predictor of exploding. (Note that those who, for example, live in majority Muslim nations will be similarly untempted by (4).)

In short, Leslie takes there to be two conditions that bring about our ready acceptance of striking property generics about Fs. We must think of Fs as an essentialised kind, and we must think that being an F is a *good predictor* of having the striking property attributed by the generic. When these conditions are met, striking property generics are readily accepted. This is why, according to Leslie and Haslanger, they are particularly pernicious for perpetuating false beliefs and discriminatory behavior regarding social groups that are stigmatized as dangerous. Each author presents slightly different stories about this, which I will take in turn.

2. Haslanger

The striking property generic claim on which Haslanger focuses most of her attention is (5):

(5) Blacks are violent.

She notes (following Leslie) that because not many instances are required for the truth of a striking property generic claim, people may accept this as true very readily. Haslanger herself wants to reserve judgment on the truth of (5), largely because she does not want to commit herself to a particular theory of the truth conditions for striking property generics. Instead, Haslanger notes that (5) is *either* false or true but very misleading.² (Leslie holds that it is definitely false because blacks are not disposed to be violent “given half a chance”, which is what she requires.³)

However, Haslanger argues that *even if* (5) is true, it should be denied. This seems paradoxical, but Haslanger argues for the use of metalinguistic negation—a mechanism first discussed by Larry Horn (1985) which allows for the denial of true claims that are in some way badly put, such as in (6).

(6) Fido didn't shit on the rug; he had a little accident.

One use to which we often put metalinguistic negation is the denial of true statements with false implicatures.

(7) A: What do you think of Z's philosophical abilities?

B: Well, he has very nice handwriting...

C: No, B, that's wrong—he's an excellent philosopher!

Haslanger suggests that even if claims like (5) are true, they should be denied using metalinguistic negation, because they carry false and pernicious implicatures. In particular, she takes claims like these to carry implicatures about natures — that there is something about the nature of blacks that makes them prone to violence. This is because striking property generics are, as we have seen, linked to essentialising beliefs about kinds. These sorts of beliefs about natures are a crucial part of the ideologies of racism, and communicating them without openly expressing them is a particularly pernicious way of propping up and disseminating these ideologies. Haslanger argues that if implicatures are not challenged they will be added to the

² The reason that she grants that (5) might be true, is because there might be a good explanation of black violence as a response to racist oppression, and this might supply a non-accidental connection that is sufficient to make (5) true.

³ This is actually quite problematic for her when it comes to (1), which she wants to maintain is true. Mosquitoes are not disposed to carry West Nile Virus “given half a chance”—they need to be exposed to it. In fact, I share Rachel Sterken's worries about whether Leslie has offered plausible truth conditions for Striking Property Generics (Sterken 2015 a, b). But here my focus is not on their truth conditions, but on whether they are particularly politically worrying.

common ground of the conversation—the background assumptions that are taken for granted. Because these are in the background, they become especially difficult to draw attention to and challenge. Since the implicature of a claim like (5) is not only false but an extremely pernicious belief to add to the common ground, it must be challenged. Haslanger recommends doing so with metalinguistic negation, as in the responses below⁴:

- No, blacks aren't violent! *People* are violent when they're placed in oppressive circumstances.
- No, blacks aren't violent! They don't share some common nature, but in fact are as diverse as any other group of people.

There is much that is right in this story, I think, but also much that is slightly off.

A first, somewhat mistaken, worry is this: Haslanger's focus is very much on metalinguistic negation rather than ordinary negation, even though she is not by any means committed to the truth of (5). One might expect her to say that if (5) is false, one should simply use ordinary negation. But a key part of her point is that we're not very good at thinking about the truth conditions for generics. She herself is not sure what they are, and— perhaps even more importantly—we are sure to encounter many people who are convinced that (5) is true. Given all this, it may well be a good idea to use metalinguistic negation in order to block the addition of the implicated claim to the common ground. This way we can protest the ideology, while sidestepping a discussion of the truth conditions of generics.

However, there are still some problems. Chief among them is the fact that it seems unlikely that a claim about natures is conversationally implicated. In order for a claim to be conversationally implicated, it needs to be the case that one cannot make sense of the utterance as cooperative without assuming the claim to be believed by the speaker. But there will be very few contexts in which one needs to assume that the speaker has beliefs about the *nature* of black people in order to understand (5) as cooperative. After all, as Haslanger notes, we don't actually talk or think explicitly about natures very often.

Nonetheless, another conversational implicature is plausible, one about relevance. If a speaker utters (5), they would in fact be behaving very uncooperatively if blackness was irrelevant to violence. To understand the utterer of (5) as cooperative, we do need to assume that blackness

⁴ Haslanger doesn't actually suggest particular responses. These are my thoughts on how the metalinguistic negation might work.

is relevant to what they're saying, an implicature which does not include anything about natures. (Compare: We'd easily judge an utterance of "People with tonsils are violent" as uncooperative, unless it turned out that there was a link between tonsils and violence.)

But once we have allowed that this is the more plausible implicature story, generic constructions cease to be a part of that story. Any time race (or some other trait) is mentioned, there is arguably just the same sort of implicature that the trait in question is believed by the speaker to be relevant—after all, it would be uncooperative to mention race if it wasn't thought to be relevant. This was strikingly demonstrated for me by a local shopkeeper, who spotted my "This is what a feminist looks like" bag, and decided to tell me some of her supposedly feminist beliefs. She proceeded to utter (8) and (9).

(8) Many Asian men are abusive to women.

(9) My friend has a friend who's Asian and he's abusive to his wife.

I found these utterances very disturbing: it's undoubtedly true that many Asian men are abusive to women, as are many men from any ethnic group; and it may well be true that a friend of her friend is an Asian who is abusive to his wife. However, there was no reason to introduce ethnicity into the discussion. Her mention of *Asian* men suggested that somehow their being Asian was relevant to the abusiveness. So, just as in Haslanger's example, there would seem to be the problematic implicature that being Asian is relevant to the dangerous property being attributed.

Importantly, though, this example is one in which no generics are used: (8) and (9) do not include generics, and yet they carry the same sort of relevance implicature that is carried by (5). Almost any time a race or ethnicity is mentioned, there will be an implicature that it's relevant to the subject matter under discussion. More generally, if specific groups of people are introduced into a discussion there will be an implicature that these specific groups are relevant to the discussion. (There are exceptions, of course. Imagine a conversation among a team of social scientists hoping to learn about differences between groups with respect to some property P. They would report their findings about these groups and property P without any implicature that there was a relationship, since this is precisely what they are trying to find out.) Introducing racial, ethnic, and others groups into discussions in this way undoubtedly makes a contribution to the transmission of pernicious ideologies. However, Haslanger has not shown that generics play a particularly worrying role in this phenomenon.

3. Leslie

Sarah-Jane Leslie raises three concerns about generics that are importantly linked. The first concern is that our tendency to overly-easily accept striking property generic claims fuels hatred and violence against members of stigmatized social groups. The second is that the use of generics to discuss social groups leads people to be more likely to essentialise social groups. This matters, because essentialising social groups makes it far more likely that we will accept striking property generic claims about these groups. Her third worry is that use of *labels* leads to the essentialising of social groups.

Leslie's proposed remedy for these problems is to alter our linguistic practices: "The evidence strongly suggests that the use of labels and generics contribute to essentialization, and so we may expect that the converse will also hold: reducing the use of labels and generics for racial, ethnic, and religious groups may reduce the extent to which children grow up essentializing these groups." (Leslie Forthcoming: 42)⁵

I will be raising concerns about both her evidence that generics are pernicious, and her proposed remedy. Along the way, we will see that some the evidence she cites actually counts *against* her proposed remedy.

3.1 Over-easy generalisations

Leslie makes dramatic use of real-world case studies to show the perils of striking property generics. These are real-world historical examples in which people reasoned from a very few instances of bad behavior to condemnation of whole groups of people. Leslie uses these case studies to argue that the human mind has a powerful tendency to make such inferences when social groups are essentialised. She writes:

Since we are working under the hypothesis that generics give voice to psychologically primitive generalizations, this observation implies that our basic way of dealing with dangerous or harmful information involves the rapid generalization of this information to the appropriate kind or

⁵ It is worth noting that in Wodak, Leslie and Rhodes (2015) this position is clarified a bit (the forthcoming paper cited in the text was actually written earlier). Here it is noted that the position is not one of total abstinence regarding use of generics for social categories, as claims like "doctors wear scrubs" is unproblematic. Leslie and her co-authors also note that abstinence would be insufficient, since non-generic statements are often recalled as generics. There is no suggestion, however, of any backing away from the thought that we should avoid the use of generics for social groups like racial and religious ones, which tend to be targets of prejudice.

category. (Forthcoming: 4)

The cases she discusses serve as powerful illustrations of this deeply pernicious tendency. They include anti-Muslim prejudice in the aftermath of September 11th; anti-Algerian prejudice in France in the 1920s, triggered by a single incident; and the way that English travelogues fuelled prejudice against both Africans and Native Americans. With their tales of racially motivated massacres and systematic discrimination, they show that this tendency to generalize from dangerous instances is of far more than merely intellectual interest. This is a vital matter of real-world political significance. If use of generic sentences plays a role in bringing about these phenomena, and changing our language could make a difference, we certainly should work to change our language. Indeed, one comes away from Leslie's paper with a very hopeful feeling—that she has identified a particularly pernicious form of speech, and that changing our speech could aid us enormously in the fight against prejudice.

However, if this form of reasoning is, as Leslie repeatedly suggests it is, simply the way that our minds deal with information about danger, it is odd to focus so much on a particular form of words. As Leslie herself says:

“To recap, then, our most primitive generalizations, voiced in language as generics, are especially sensitive to information that is particularly striking, horrific, or appalling. When we encounter individuals engaging in such an act, we are naturally inclined to seek to generalize this action to a kind to which the individuals belongs” (Leslie Forthcoming: 22)

Importantly, Leslie's claim here is that *witnessing a single horrific incident* will cause us to engage in pernicious generalizations. The cause of the pernicious generalization is simply knowledge of the single incident (combined, presumably with a tendency to essentialise the relevant kind). The generic language comes after the fact—it allows us to express the generalization that we have already made.

Even Leslie's own examples do not convincingly show that striking property *generic* statements play any key role in the horrific prejudices she discusses. Indeed, some of her own examples are of other sorts of utterances, like Saxby Chambliss suggesting that the sheriff should “arrest every Muslim that crosses the state line” (Leslie Forthcoming: 11) a clear universal generalization. If Leslie is right that our minds tend to leap by default from single dangerous instances to very general beliefs, then we should expect this to happen regardless of the words that are used. And in fact she describes just this phenomenon in discussing the way that English colonial narratives fuelled prejudice against Africans: “it would be overly simplistic to place the

blame for the formation of early negative stereotypes squarely on the explorers. Many of them were quite responsible in their reporting, and did not indulge themselves in broad generalizations. Given the nature of our default system of generalization, they did not have to. The reporting of specific instances would suffice to encourage very general beliefs in the mind of the reader” (Leslie forthcoming: 13-14).

The stereotypical beliefs formed by the reader were, as Leslie herself notes, formed on the basis of individual instances reported. I think Leslie is absolutely right about the mind’s natural tendency to generalize from a few dangerous incidents. But if this is right, then it’s not at all clear that avoidance of generic utterances would help to counteract this. We will remain just as likely to make these pernicious inferences about social groups, whether or not we hear the groups described with generic language. Reasoning from shocking incidents to sweeping generalisations is, as Leslie argues, simply what we do.

In fact, there are two strands of thought in Leslie that are somewhat in tension. One strand focuses on our thought, and on our innate tendency to make pernicious generalisations about social groups on the basis of single incidents. The other strand focuses on the language in which we express such generalisations, suggesting that a reform of our language could reduce our tendency to perniciously generalize. But if the first is right, then the second seems very unlikely: if merely encountering single incidents causes us to generalize, then avoiding particular phrases by which we might express these generalisations will not be an effective remedy.

3.2 Generics increase likelihood of essentialising

Leslie makes use of several studies which purport to show that use of generics increase the likelihood that subjects—both children and adults—will essentialise a category. (See for example Cimpian and Markman 2009, 2011; Gelman, Wear and Kleinberg 2010.) It is beyond the scope of this paper to describe each of these studies in detail. However, one common feature that they share is that they contrast the effects of generic sentences (*Fs have property P*) with non-generic sentences (*This F has property P*; *This has property P*). These studies show that after hearing/reading a generic sentence like *Fs have property P*, subjects are more likely to essentialise Fs than after hearing a non-generic sentence like *This F has property P* or *This has property P* (where it is clear that ‘this’ picks out an F). This is shown through several measures of essentialising, for example: a tendency to explain the presence of property P in terms of nature, function or stable trait; a tendency to infer that things with property P are likely to be F; a tendency to infer that Fs will share other properties besides P; “a tendency to expect properties associated with the category to be innate or inevitable” (Rhodes et. al. 2012: 2). Work that Leslie

herself was involved with shows that these findings extend to social, not just biological categories.

However, these studies share a common flaw. They contrast *Fs* with *This F* or *This*.⁶ A claim that *Fs have property P* attributes a shared property to a greater number of *Fs* than a claim that a particular *F* has property *P*. If all that I know about *Fs* is that they share some property then it will seem (at least somewhat) reasonable to infer that *Fs* have a shared nature.⁷ By contrast, if all that I know about *Fs* is that one of them has some property, I'll be more hesitant to attribute a shared nature to *Fs*. But there is no good reason to believe that *generics* are playing a key role here. It is entirely possible, given the evidence, that other attributions of a shared property would have exactly the same effect. Here are several alternative shared property attributions which might have the same effect, but which have not been tested:

- Many *Fs* have property *P*
- Most *Fs* have property *P*
- This *F* has property *P*. And this *F* has property *P*. And this *F* has property *P*. And...⁸

If my supposition here is correct, then these property attributions may also lead to essentialising. And yet, they have not been tested.⁹ All that we know so far is that attributing a property to *Fs* is more likely to lead to essentialising than attributing a property to *an F*. It seems to me that this is an insufficient evidence base for inferring a special problem with *generics* and trying to re-phrase our utterances.

But there is also a further problem with using *This F has property P* as our contrast item. I suspect that this form of words will, at least in contexts where we're being introduced to *Fs*, carry the implicature that *Fs* with property *P* are not the norm. That is why it would be very strange for a children's book to introduce cats with a sentence like "this cat has a tail", or for me to tell you that people on my street have heads. It's not at all strange for a children's book to say "this cat has black spots", and I think this is because (while they're not especially unusual), such cats are not the norm. Now, of course the presence of an implicature like this doesn't make you doubt

⁶ An exception to this is Gelman et. al. 2002, who contrast the generic with 'all' and 'some', finding that are children and adults are most likely to make inductive inference from 'all' statements, less likely with generic statements, and less likely still with 'some' statements. This study does not test the alternatives suggested above.

⁷ This is especially likely if it is an unusual property, as it is in many of these studies (e.g. being afraid of ladybugs; disliking ice cream).

⁸ Thanks to Elena Hoicka for suggesting this alternative.

⁹ I am currently working with Elena Hoicka and Rachel Sterken to test these.

that a kind you already know about is essentialised—you'd still think cats have a nature. But if you introduce a new kind only by listing properties in a way that suggests they're not typical, that might well give one pause in inferring to the existence of a nature. (Contrast “this gluk likes tea” with “most gluks like tea”.) After all, one would normally introduce a new kind by discussing typical properties. If one is instead listing off atypical properties, it would be reasonable to wonder if there are typical ones. This suggests, then, that *This F has property P* is an especially bad choice of contrast term for the generic claim.

The objections raised here are not meant to defeat Leslie's claims. All I have done is to raise some possibilities that I think need to be investigated before we can conclude that generics have the very particular role in essentialising that Leslie takes them to have. It seems to me that, for all we presently know, it may be a much broader class of constructions that play this role. This, then, is merely a demand for more research before we embrace Leslie's conclusion.

3.3 Labels

Leslie also makes use of studies on the use of *labels*. These are studies (e.g. Gelman and Heyman) that show children are more likely to do things which are closely related to essentialising when labels are used to describe someone. For example, if children are told “Rosie is a carrot-eater” they will take the carrot-eating to be a far more stable property of Rosie's than if they are told “Rosie eats carrots whenever she can”. Since essentialising requires taking a kind to have a shared and stable nature, Leslie suggests that use of labels may promote essentialising, and that avoidance of them in discussing social groups might help to reduce essentialising. She finds further evidence of this in work by Waxman (2010), which shows that preschoolers presented with images of members of different racial groups are more likely to assign significance to racial categories if the people are labeled (e.g. as ‘a Wayshan’) than if they are merely described (e.g. as liking to “play a game called ‘Zaggit’”).

But if the evidence above is right, it is again a red herring to focus on use of generics—the problem is with labels for social groups. Granted, use of generics requires use of noun phrases (which labels are), but avoidance of generics will not succeed in bringing about avoidance of noun phrases for social groups. So it will give a false sense that one is helping to solve the problem, when one is not. Admittedly, Leslie argues for *both* avoiding social group generics and avoiding social group noun phrases. But by avoiding social group noun phrases, we'll be avoiding social group generics anyway. So discussing generics in this context seems a distraction.

4. How racism is transmitted

An unarticulated assumption underlying Leslie and Haslanger's work is that racism is, to some significant extent, transmitted via utterances which make explicit reference to racial kinds, and attribute a dangerous property to members of these kinds. There is, of course, no doubt that such claims have been, and still are, made. At some times and in some places, these utterances may play a major role in the transmission of racist ideology.

However, much racist ideology is transmitted without explicit mention of race. Some of this is linguistic, as with the use of dogwhistle terms (Saul forthcoming, Stanley 2015). But a good deal of it will not be—a good deal of ideology is transmitted simply by living in a world structured by residential, occupational, and educational segregation (Anderson 2010). Crucially, this includes the transmission of racism to children. Indeed, many white parents believe that they should never mention race to their children (Vittrup and Holden 2011). These parents would never dream of using sentences like (5) in conversation with their children. Leslie's suggestion is that by changing the language that we use to teach children about racial groups, we can fight racism. But now consider more carefully the nature of parent-child discussions of race. There are undoubtedly white explicitly racist parents who explicitly teach their children about race. However, these parents are not those well-meaning liberal parents for whom Leslie's message is intended. Among well-meaning white liberal parents the problem is not the use of generics in discussing race. Instead it is more likely to be the absence of race discussions, which could lead children to question the racially-structured world that they live in. Telling such parents to avoid the use of generics is unnecessary, and may only shore up their belief that discussions of race are so tricky that it would be better to avoid them.¹⁰

5. The Benefits of Generic Language

Although Leslie and Haslanger's focus is on the ways that generics can serve to perpetuate racism, sexism and other ills, it also seems to me worth noting the ways that they can serve to *fight* these ills. Here are some generic claims that campaigners for social justice might well want to make, as part of a social critique:

(10) Women are expected to want children.

(11) Black people face discrimination.

¹⁰ Black parents are far more likely to explicitly discuss race with their children, in order to make them aware of racism. But our focus here is on how the racism of white people is perpetuated.

- (12) Gay people are subjected to violence.
- (13) Muslims are profiled by airport security.

If it's important to avoid the use of generics about social groups, then it's important to avoid these sorts of statements too. If use of generics leads to essentialising then we shouldn't be using them *at all*. In fact, if Leslie is right about this, we should be especially wary of using them with children. So we shouldn't say things like:

- (14) Boys like pink too.
- (15) Girls play football.
- (16) Muslims are celebrating Eid, so wish your friend Ali a happy Eid!

But these are quite clearly very good things to say to children. And, according to Leslie, they make use of a form of generalization that comes especially easily and naturally to children. Communicating with children and teaching them about the world is actually very challenging at times—especially when we are seeking to combat the prejudices that they see acted out in the world all around them. Depriving us of statements like (14)-(16) is depriving us of some very important weapons in our anti-prejudice arsenal. We may leap overly easily to reject generics entirely when we focus just on the nasty ones. We need to remember the good ones as well.

6. Risks of Generic—and other-- Language

6.1 *Yes, there is a problem*

Despite all that I have argued above, however, I do think that use of generic language carries risks. These risks result, it seems to me, from a property of generics that both Leslie and Haslanger mention, but which is not the main focus for either of them. This is the difficulty that we have in thinking clearly about them. Both Haslanger and Leslie note that while we accept striking property generic statements on the basis of just a few instances, we make universal inferences very quickly on the basis of this acceptance. I would go further than this and note many ways in which we shift around in our interpretation of these claims—ways that render us very vulnerable to manipulation (and self-deception) by way of generics.

It is very common, for example, for people to reassure themselves that they are not racist against black people by noting that racists dislike black people, and then remarking that they have a black friend. This involves starting from a generic claim which may seem reasonable (though it is, I think, based in an overly restrictive understanding of 'racist'), and moving to an

interpretation of this claim as a universal generalization. This last move is why it may seem—wrongly—that citing a single black friend shows it to be false that one dislikes black people. More generally, it is too easy rhetorically to deny a true generalization made with a generic by citing a single counter-instance.

6.2 But what is the correct solution?

But it does not follow from this that the correct remedy is to avoid use of generics. They are useful in the cause of social justice, as I noted in the previous section. But, perhaps more importantly, we won't really be able to avoid them if they do in fact "give voice to" a default mode of generalization for our minds. Leslie in fact suggests that we should try to eschew 'labels' and opt instead for 'descriptions'. And it might be true that when we initially replace 'Muslim' with 'person who follows Islam' (a suggestion that she makes) we'll be slower to ascribe an essence. But soon that phrase will simply be a label, and function as one.

If generics are how we express a default mode of reasoning, we will start making generic claims with 'person who follows Islam'. There is no technical difficulty in constructing sentences like 'People who follow Islam are terrorists'. And this is not just a speculative point. It is worth reflecting also on how notably unsuccessful it was to replace the noun 'moron' with the descriptive phrase 'mentally retarded person'. The more recent terms 'special needs' and 'person with special needs' also provide a revealing case study. Indeed, the drive to label groups with noun phrases has led the noun phrase 'special needs' to be used as an adjective in 'special needs children'. Even when it's (initially) ungrammatical, we will find a way to form the easy noun phrases that facilitate essentialising. Reflecting on cases like these should give one pause about the efficacy of attempting to reduce essentialising through this sort of linguistic reform.

So, it seems to me, we *shouldn't* try to avoid generics. Instead, we need to get better at talking and thinking about them. We need to press people to spell out their evidence for their generic claims and to reflect on what that evidence really does or doesn't warrant. Suppose that A asserts 'Muslims are terrorists'. B could ask any of the following.

- A few Muslims are terrorists. But so are a few Christians. Remember Timothy McVeigh?
- Why do you think that? How many terrorist Muslims do you know about out of the more than a billion in the world?

- Do you have any evidence that being Muslim makes you more likely to be a terrorist?

Doing this is similar to what Haslanger suggests, and indeed one could even use metalinguistic negation:

- No, Muslims are not terrorists. A few terrorists who are members of the group aren't enough to justify what you said.

However, the strategy is not confined to metalinguistic negation, and it's not motivated by the thought that there is an implicature about natures in need of cancellation.

6.3 Applying the solution more broadly

Importantly, we need to be equally concerned about other, non-generic gratuitous uses of kind terms. If an act of terrorism was explicitly carried out in the name of Islam, then there's no doubt that the perpetrator being Muslim is relevant to what's being reported, and so noting this fact should not be criticized. (Of course, it would nonetheless be very important to avoid overgeneralizations, whether generic in form or otherwise.) However, social kind membership is often mentioned when it's not relevant. This has been documented especially well with respect to race and crime (Anderson 2010:56). This also makes us vulnerable to manipulation and plays an important role in perpetuating prejudice. We all too easily slide—as Leslie would predict—from these individual instances to beliefs about the group, whether we use generics like 'blacks' to express the beliefs or non-generics like 'many blacks' or 'most blacks'.

Importantly, though we may not even form beliefs (consciously or unconsciously) in the relevant propositions. Instead, we may well form implicit biases. On most theories, these are not beliefs, but something more like patterns of associations, which nonetheless lead us to biased judgments. A classic case of this would be the well-documented Weapon Bias: we are more likely to judge an ambiguous object to be a gun when it appears in the hand of a black man than a white man; these biased judgments correlate with our levels of implicit bias against black men (Payne 2006). Implicit bias is largely an unconscious phenomenon, and it is a crucial part (though only a part) of the prejudice we see in action in the world. Blocking certain means of expressing prejudice—via generics such as “black men are violent” would not do much at all to combat implicit bias.¹¹ Many of those who harbour these associations are unaware of them, and

¹¹ I don't mean to suggest that Leslie takes her solution to be a cure-all for prejudice. She does not. However, the importance of implicit bias in our world enhances my worry that avoiding generics is a misdirection of our energy.

therefore not likely to seek out ways of expressing them. So rather than focus on how biases are expressed, we need to fight the biases.

However, the utterances that we are exposed to probably do make a difference to the formation and perpetuation of these biases. What matters about these utterances, though, is not likely to be whether they contain generics or not. What matters is, for example, the tendency to mention the race of criminals when the criminals are black, and not when they are white. There are at least two ways in which this could contribute to the formation of implicit biases. One—less likely—would be via the implicature mechanism that I discussed earlier: *the reporter has mentioned that the mugger was black. They must think this is relevant. They know their facts, so blackness must be connected to the crime.* But this is unlikely.¹² There's nothing remarkable about race being mentioned in this way in a crime report (indeed it is the norm to mention race of perpetrator when they are black), so we won't feel a need to reason about why it would be happening. What's far more likely to happen is that mentioning race when criminals are black leads us to associate blackness with criminality—*whether we formulate this as a belief or not.* (Leslie would predict that just a very few mentions would suffice for this effect, but in fact we will have been exposed to far more than a few such instances.) And this association, we know from decades of work on implicit bias discussed earlier, is far from benign.

In order to take action on this, we need to start speaking up about these gratuitous mentions of race:

- Why did they mention that mugger was black? Is that meant to be relevant?
- Why should we care about the mugger's race?
- Why don't we ever hear mentions of race when white people are criminals?

My view, then, is that we would be making a mistake to focus our efforts on *generics*. Instead, we need to think hard about all mentions of gender, race and the like. We need to call attention to them and ask whether they're relevant, demand evidence for claims, question inferences, and so on. We definitely *should* do these things when generics are used. But we should also do them

¹² Of course, an implicature can arise without the calculation actually taking place. However, the calculation has to be part of a rational reconstruction of how the belief that blackness is connected to crime is arrived at. That is, it has to be that taking the speaker to believe this is required to make sense of her utterance as cooperative— even if the audience doesn't reason explicitly about what's required to understand her as cooperative. And it's just not true that this is required to understand the speaker as cooperative—mentioning blackness of criminals when reporting on crime is such a norm that it's simply following the conventions of the genre.

when generics are not used.¹³

Anderson, Elizabeth. 2010. *The Imperative of Integration*.

Cimpian A, Markman EM (2009) Information learned from generic language becomes central to children's biological concepts: Evidence from their open-ended explanations. *Cognition* 113:14–25.

Cimpian A, Markman EM (2011) The generic/nongeneric distinction influences how children interpret new information about social others. *Child Development* 82:471–492.

Gelman, Susan. *The Essential Child* (New York: Oxford University Press, 2003).

Gelman, Susan; Wear, Elizabeth; Kleinberg, Felicity. 2010. "Effects of Generic Language on Category Content and Structure." *Cognitive Psychology* 61 (3): 273-301.

Gelman SA, Star J, Flukes J. 2002. "Children's use of generics in inductive inferences." *Journal of Cognition and Development*. 2002; 3:179–199.

Haslanger, Sally. 2011. "Ideology, Generics, and Common Ground" in Witt, Charlotte (ed.), *Feminist Metaphysics*. Springer Verlag 179-207.

Horn, Laurence. 1985. "Metalinguistic Negation and Pragmatic Ambiguity". *Language* 61: 1, 121-174.

Leslie, Sarah-Jane. Forthcoming. "The Original Sin of Cognition: Fear, Prejudice, and Generalization." *Journal of Philosophy*.

Mendelberg, Tali. 2001. *The Race Card: Campaign Strategy, Implicit Messages, and the Norm of Equality*. (Princeton: Princeton University Press.)

¹³ I have benefited from many incredibly helpful discussions of this work, especially with Elena Hoicka and Rachel Sterken, but also with a wide variety of lovely philosophers at Stirling, St Andrews (Workshop on Haslanger's Ameliorative Project), Oslo (Workshop on Conceptual Engineering), Waterloo, Barcelona (GRSelona 2 Conference), Harvard (The Generic Workshop), Vermont, and Tartu.

Payne, K. 2006. "Weapon Bias: Split-Second Decisions and Unintended Stereotyping", *Current Directions in Psychological Science* 15: 6, 287-291.

Rhodes, Marjorie; Leslie, Sarah-Jane; and Tworek, Christina. 2012. "Cultural Transmission of Social Essentialism." *PNAS* 109 (34):13526-13531.

Saul, Jennifer. Forthcoming. "Dogwhistles, Political Manipulation, and the Philosophy of Language". Forthcoming in Daniel Fogal, Matt Cross and Daniel Harris's *New Work on Speech Acts* volume (under consideration with OUP).

Stanley, Jason. 2015. *How Propaganda Works*. Princeton: Princeton University Press.

Sterken, Rachel. 2015a. "Leslie on Generics". *Philosophical Studies* 172 (9): 2493-2512.

Sterken, Rachel. 2015b. "Generics, Content and Cognitive Bias". *Analytic Philosophy* 56 (1): 75-93.

Vittrup, Brigitte ; Holden, George W. 2011 "Exploring the Impact of Educational Television and Parent-Child Discussions on Children's Racial Attitudes", *Analyses of Social Issues and Public Policy*, Vol.11(1), pp.82-104

Waxman SR. 2010. Names will never hurt me? Naming and the development of racial and gender categories in preschool-aged children. *Eur J Soc Psychol* 40:593- 610.

Wodak, Daniel; Leslie, Sarah-Jane; and Rhodes, Marjorie. 2015. "What a Loaded Generalization: Generics and Social Cognition". *Philosophy Compass* 10 (9): 625-634.