

Social and Political Aspects of Generic Language and Speech

Matthew McKeever & Rachel Sterken

This is a penultimate draft, and may differ in tiny ways from the published version.

We often generalize about the world, including the social world, using *generics*. We say that ducks lay eggs, that nurses can diagnose covid-19, that bonds are safer than stocks. Being able to say such things, and think the corresponding thoughts, is important for us. In making sense of ourselves and the world, we often appeal to generalizations and the underlying explanations which support them. If, for example, one moves to Hong Kong, it's very helpful to know that Hong Kongers speak Cantonese and drive on the left side of the road. There are of course exceptions, but these kinds of generic generalizations¹ – generalizations like *Hong Kongers speak Cantonese* and *nurses can diagnose covid-19* – are distinctively useful in many settings. If you're a nurse, it can be useful to know that nurses are overworked and underpaid. That helps you understand yourself and to act on that information in various ways. If you're a child, it's good to know that in many jurisdictions, children are protected by law from abuse. In these and indefinitely many other cases, generic claims help us to understand ourselves, the world, our place in the world, and to act on such knowledge.

But generics are also potentially hazardous. Recent work by philosophers, linguists and psychologists have pointed to numerous features of generics which suggests that, semantically, cognitively and epistemologically, generics are complicated, and that this can have bad effects. This chapter will explore these complexities, and in due course we will review some of the relevant literature, but let's present one such issue quickly to give a sense of what can potentially go wrong. Consider:

(1) Immigrants commit crimes at a higher rate than native citizens.

Even if we assume that (1) is false for most (or maybe all) countries, it is a claim that many people believe. That claim is often ruthlessly exploited. Claims like (1) play important roles in political manipulation. Recent work in philosophy and psychology suggests that it is distinctive features of generics that makes them exploitable in this way. For one illustration, note that (1) doesn't say *how many* immigrants commit crimes. As we'll see, some generic sentences of the form *Fs G* require many *Fs* to *G* for their truth, while others require few: generics are semantically variable in peculiar ways. A politician can exploit this fact. Even if 98% of the immigrant population is law abiding, if there are some high-profile crimes committed by immigrants, the politician can point to them and say: "look, I told you, immigrants commit crimes," and get away with it. That can strike us as an acceptable move in the political language game, even if it is specious. (Note the parallel reasoning with a universal quantification would be an obvious non-sequitur.)

So the semantic variability of generics is a tool bad actors can exploit. Generics are also cognitively tricky. Again to take just one example to be discussed more later, experiments suggest that people tend to

¹ The literature on generics is vast. Good places to start exploring it include Krifka et al. (1995), Leslie and Lerner (2016), Pelletier (2010) and Sterken (2017).

accept a sentence like (1) on the basis of little evidence, but, having accepted it, come to think that it makes strong claims on reality. A study conducted by Andrei Cimpian, Amanda Brandone, and Susan Gelman (2010) found that, on average, upon hearing that 69% of a novel kind possessed a given property (e.g., 69% of Zarpies eat flowers), subjects tended to accept the corresponding generics (e.g., *Zarpies eat flowers*) as true. On the other hand, again on average, upon hearing the same generic, subjects expected 96% of the novel kind to possess the property. Worse still, if the predicated property happens to be striking or appalling, as in the much discussed example:

(2) Mosquitoes carry West Nile virus.

The former percentage may be much lower (Leslie 2017; for some relevant empirical work, see Abelson and Kanouse 1966 and Khelmani, Glucksberg and Leslie 2011).

These cognitive mishaps provide a further explanation for why sentences like (1) are particularly hazardous. A person can come to accept (1) having read about a few instances of criminality by immigrants (which will of course exist, because there are criminal immigrants just as there are criminals in any group) but then come to inflate it and think that there are many immigrant criminals, and thus, that restriction on immigration is a worthy political goal, and that the anti-immigrant politician should be elected.

Faced with this and a host of other troubling observations about generic representation, some philosophers and psychologists have asked: shouldn't we stop using generics?

The goal of this chapter is to present a plausible case for an affirmative answer, to present some extant reasons for preferring a negative answer, and then to present some further reasons for thinking that, their flaws notwithstanding, we shouldn't get rid of generics, rather we should think about how we can improve them.

Section 1: The why of avoiding generics

Above we introduced the idea that speakers should limit their use of generics and gave a few brief examples to illustrate why. This section will consider more systematically the range of reasons that might motivate avoiding generics.

1.1 Generics as semantically defective

There are a variety of reasons one might have to prohibit or avoid (the use of) an expression. For one, an expression might be *semantically defective*. This is a familiar point from the recent literature on

conceptual engineering:² representational devices can be defective in many ways, and once we realize this, we should seek to develop better representational devices and get rid of inferior ones.

Let's take a small selection of the ways we know representations fail and see if generics fit the pattern. You might think, simply enough, that generics as a class or the generic operator *Gen*³ are like *phlogiston* or *the golden mountain*⁴-- empty terms. Indeed, we think that something roughly like that can be found in the literature that prescribes that speakers avoid generics: If one thinks that generics purport to refer to kinds and their essences, and one thinks that there are no essences for large swathes of the social world, then racial and gender generics might just simply be expressions of an out-dated theory. That's a live possibility, and one that we wouldn't want to dismiss out of hand. Indeed, Hom (2018) has a theory according to which slurs are semantically empty and are deficient in meaning, because their truth-conditions depend on the existence of essences which determine, or would determine, the application of a given slur. There's a case to be made for saying that racial or gender generics are something like Homian slurs, which is to say *empty* (and *slurring*) because *essentializing*.

Another approach would have it that generics are *tonk*-like. Recall *tonk*, Arthur Prior's (1960) example of a connective whose introduction and elimination rules are out of harmony. From *A* one can conclude *A tonk B* and from *A tonk B* one can conclude *B*: from a sentence one can infer a conjunction of that sentence with any other sentence, and from any conjunction one can infer any of its conjuncts. For arbitrary *A* and *B*, *tonk* lets you reason from one to the other, which is to say a logic with *tonk* is no logic at all.

The reason why this deserves a bit of thought is because, as Leslie has noted in passing (2017:397; see also Cappelen and Dever 2019, which discusses the idea as well), it does seem like (something akin to) the introduction and elimination rules of generics don't match. The experiments we cited above showed that we introduce generics on limited evidence but eliminate them for very strong claims on reality, and on one slightly rough way of looking at it, that's what we do with *tonk*.

That being said, to make the case on that basis that generics are *semantically* defective is not yet well motivated (not every mistake we make with language is owing to semantics). More work is needed to establish that *Gen* is semantically defective. It could be that *we* are defective in our handling with generics, a possibility which we'll consider now.

1.2 Generics lead to cognitive and epistemic errors: amplification, degrading, shifting and entrenchment

² See Eklund, this volume, for an introduction to conceptual engineering.

³ A covert operator that expresses the peculiar quantificational force of generics. See the references in the first footnote for more on *Gen*, about which we won't say more here.

⁴ In a framework that requires definites to refer to be meaningful.

One of the interesting and distinguishing features of generics is the extent to which their use leads us into epistemic and cognitive errors and confusion. A simple argument then suggests itself: if an expression leads us into epistemic and cognitive errors and confusion, we shouldn't use it.⁵

In this section, we'll consider four forms of mistakes and confusions that generics seem to lead to more easily and detrimentally than other expressions. We'll call them *amplification*, *degradation*, *shifting*, and *entrenchment*. As we've set up some of the discussion below, we're presenting these as mishaps at the level of thought, which could take the form of mishaps of reasoning, cognitive negligence, willful ignorance, and information retrieval, among other things. It's worth noting that many of these could be mishaps at the level of communication and interpretation as well, or even at the social epistemological level.⁶ It's a largely open question just where these mishaps are happening; potentially they're happening everywhere. Precise formulations of amplification, degradation, shifting and entrenchment will obviously differ from what's presented below depending on the locus of the misap. There may be cross-cutting and interaction of these categories that we're glossing over.

Amplification

Here is a statement of what amplification might consist in:

Amplification: Subjects have a tendency to accept generic claims based on relatively weak evidence, but then later take generic claims to have rather strong conclusions.

There are three important types of amplification error that in one form or another can be found in the literature, and we'll go through each in turn: statistical, modal and normative.

Statistical amplification.

A good place to start is with one of the canonical sentences in the literature, that we saw in the introduction, repeated here:

(2) Mosquitos carry West Nile virus.

The facts are that fewer than 1% of mosquitoes carry West Nile virus, but the generic is nevertheless completely fine. Leslie, who calls such generics 'striking property generalizations,' suggests that there is

⁵ One cognitive and epistemic error that we're excluding from this section is the error that arises when one uses a semantically defective term. We're assuming here that the meaning of generics is fine, but that we misuse them. The coherence of this assumption is not unquestionable (one might think a term's epistemic badness implies its semantic badness), but we won't question it here.

⁶ Compare: it might be that to use a singular *term* you must be acquainted with its bearer. Or it might be that to have a singular *thought* you must be so acquainted. Or maybe neither of these holds, but as a matter of pragmatics one *shouldn't* use a term if one isn't acquainted. These location questions are important, but we don't have space properly to address them here.

something special about them relative to other generics, and suggests that their specialness consists in the fact that:

their predicates express properties that we have a strong interest in avoiding. If even just a few members of a kind possess a property that is harmful or dangerous, then a generic that attributes that property to the kind is likely to be judged true. (2017: 396)

When we are dealing with harmful or dangerous properties, we quickly generalize on the basis of few instances to the kinds to which they (purport to) belong.

Now, in and of itself this isn't so bad. It's maybe a bit weird that generics are like this, but it might be a helpful feature that they are. The cognitive problems arise when we examine the role such striking property generalizations play in our lives.

One such problem is that, even though we *accept* a sentence such as (2) on the basis of few instances, experiments in psychology suggest that we *conclude on the basis of* such a sentence that it has many instances -- recall the Cimpian, Brandone and Gelman (2010) study summarized in the introduction which found just that.

With that said, we can now immediately see the importance of these results for socially and politically engaged philosophy of language. Leslie is much concerned with sentences like (3), repeated here:

(3) Muslims are terrorists.

It is like (2) in the respect just presented: it predicates of a group a dangerous property that few of its members are likely to have. We can then conclude that people will be wont to accept it, and thus to falsely conclude on the basis of it that many Muslims are terrorists.

Let's take another case to really push this home. Consider the following statement:

(4) Mexicans bring drugs and crime to the US. Mexicans are rapists. Some are good people.

The reader will of course know that we're closely paraphrasing Donald Trump's remarks. For the sake of propriety, let's just quote them exactly:

(5) When Mexico sends its people, they're not sending their best.... They're sending people that have lots of problems, and they're bringing those problems with us. They're bringing drugs. They're bringing crime. They're rapists.

In response to this, Tim Kaine claimed that Trump said *all* Mexicans are rapists. During a debate with Mike Pence:

Kaine: These guys say all Mexicans are bad

Mike Pence: That's absolutely false.⁷

Pence went on to mention the qualification about some being 'good people'. Kaine's remarks, in particular his universal quantification, were fact-checked by reputable sites,⁸ and found false.

Now, as far as we know, not much came of this. But note what we're seeing: a perhaps slightly careless amplification of an unclear vitriolic remark that enabled Pence, truthfully, to defend Trump. At the same time, there's no doubt that many Trump supporters (or, at the time, would-be supporters) were more than happy to amplify his generic and take him to be slandering all Mexicans.

There's a lot more to say on this theme, however we'll have to limit the rest of our discussion to a final observation and case involving inferences of exhaustivity. As we write this, the bad treatment of Black people in the US by police has once again been made salient, and a now long running debate has played out on social media and in newspapers. The protesters rally under the claim that black lives matter, and the response comes that *all* lives matter.

We can note that *black lives matter* is a generic, and 'all' is an amplifying device, so a natural question to ask is whether we can explain this rhetorical move in terms of some particular features of generics.⁹

To see how, note some basic facts about the inferences licenced by different quantifiers, familiar from generalized quantifier theory. In particular, we'll present without argument, these facts, instantiated with a simple example for the sake of illustration:

- (i) Some blue marbles roll \Rightarrow Some marbles roll
- (ii) All marbles roll \Rightarrow All blue marbles roll
- (iii) All blue marbles roll $\not\Rightarrow$ All marbles roll
- (iv) Some blue marbles roll $\not\Rightarrow$ \neg All marbles roll
- (v) All blue marbles roll $\not\Rightarrow$ \neg All marbles roll

The first three give a sense of the differing inferences licenced by different quantifiers (which we won't get into any further here: see Barwise and Cooper 1981 or Westerståhl 2019). (iv) and (v) are the important ones. Neither existential nor universal quantifications of qualified noun phrases imply the falsity of the universal unqualified noun phrase sentence (they probably implicate it; see the preceding footnote). That means to say, if the all lives matter people were interpreting the generic as existential or universal quantifiers they wouldn't, as a matter of logic, have anything to complain about. However, note:

- (vi) All, *and only*, blue marbles roll \Rightarrow \neg All marbles roll

⁷ See next footnote for a source and some background.

⁸ www.politifact.com/factchecks/2016/aug/08/tim-kaine/tim-kaine-falsely-says-trump-said-all-mexicans-are/

⁹ Others have suggested, plausibly enough, that we can understand these exchanges in terms of implicatures: <https://www.languagejones.com/blog-1/2016/7/8/the-linguistics-of-blm-scalar-implicature-and-social-controversy>.

(Assuming there are marbles other than blue ones.) Adding what we can call an exhaustivity implicature or explicature to the universal qualified noun phrase *does* get us something entailing the falsity of the universal. By parity of reasoning:

(vii) All, *and only*, black lives matter $\Rightarrow \neg$ All lives matter

It follows that if a hearer were understanding *black lives matter* as both amplified and exhaustive, they would be hearing something that, as a matter of semantics (plus the empirical fact that there are lives other than black ones), contradicts something the *all lives matter* person says, and that would give us some explanation of why some people chose to respond in such a way. Accordingly -- but this is tentative -- we think it's worth considering whether yet another feature of generics worth paying attention to in social and political settings is their tendency to be interpreted as exhaustive (for some extant work on exhaustivity and generics, see Nickel 2010; for more on the pragmatics of such statements, see Anderson 2017.)¹⁰

Modal amplification

The generalizations expressed by generics don't always simply consist of statistical information -- witness examples like:

- (6) Supreme Court Justices have even social security numbers.
- (7) OrangeCrusher2000s crush oranges.

(6) conveys a generalization that is stronger than its universal counterpart: In a scenario where all Supreme Court Justices have even social security numbers, (6) can be intuitively false. While (7) can be true even when the said juicers are directly off the assembly line and are never going to be used -- the truth of (7) seems to stem from the juicers' *capacities*, a distinctively modal notion. These cases show that sometimes generics -- harmlessly enough -- convey information about kinds or essences, or other modal content which is part of the generalisation expressed. And this might be useful: it might be good to have a bit of language that enables us to speak about the nature of lizards or bonds or madrigals, as opposed to what merely happens statistically to hold of many of them. But this fact can also lead to error. We call this *modal amplification*: modal amplification occurs when a subject accepts a generic claim on the basis of weak modal or statistical information, but then concludes from that generic that there is a much stronger modal relation which holds between the given kind and property.

¹⁰ You might object to this analysis as so: really what's going on, you might think, is a question of formal pragmatics. The black lives matter person proposes as a question under discussion 'What are some important facts about black lives?' The opponent changes the question to 'What sort of things matter?' distracting from the important and timely question posed. Phenomenologically, you might think, this *feels* more true to the case in question. Maybe so. But in our opinion, these sorts of location questions (here: whether to locate the phenomenon in semantics or formal pragmatics) are subtle, and we don't think that letting several flowers bloom, in an explanatory spirit, is a bad tack (just as we think we should heed both pragmatic and semantic accounts of domain restriction, to take but one example).

Sally Haslanger (2011) discuss cases where subjects come to accept a generic like:

(8) Women are submissive (nurturing, cooperative).

perhaps based on having witnessed several instances or based on testimony to that effect, but then come to accept that it's *in the nature of* women to be submissive (nurturing, cooperative).

In order properly to make this point, we need to briefly review some material from metaphysics. A venerable tradition in philosophy stemming from Aristotle and recently reinvigorated by Kit Fine (for example, 1994) claims that objects have essences. In the words of Haslanger, who has done important work on modal amplification, the essence of an object is "a set of intrinsic qualities that explains the characteristic behavior of things of that kind" (Haslanger 2011: 186) such that, if the object ceases to possess those qualities, it cease to be the type of thing it is.

It is very natural to understand many generics in terms of the notion of essence. A species that was duck-like but didn't lay eggs, we think, wouldn't be a duck: sometimes generics express essences, and essences, in turn, warrant counterfactuals about the existence or lack thereof of an object in other possible worlds (in a world in which Daffy reproduces asexually, Daffy isn't a duck).

But now we can see the possibility of modal error lurking. We hear, or come to believe, a sentence like (8), because the sitcoms we watch, or whatever, tell us, and it doesn't clearly jar with our experience. But if we think that generics are used to express essences, as they seem quite clearly when we consider, say, ducks, then we might think that it is of the essence of women to be nurturing.

Similarly, it seems less clear that we want to accept 'men are nurturing,' and so are less likely to think that it's of the essence of men to be nurturing. But this is a dreadful result! It's false that women are nurturing in a way that men fail to be just because they're women. It's false that if someone fails to be nurturing, she ceases to be a woman. And we don't want to give the (admittedly, atypically metaphysically sophisticated) male chauvinist the chance to say 'You should be the one to console our child--it's in your nature to do so!'

This point bears lingering on. Indeed, it's arguable that these sort of essentializing concerns are one of the main reasons that people find generics objectionable. Anderson, Haslanger, and Langton (2012) argue that:

if a characteristic racial generic is asserted, it should be rejected, because is false, and also politically problematic: it presents social artifact as racial essence (765)

There are different ways to read this, but on one interpretation, it's saying that a big problem with generics is that they lead us into modal amplification, tricking us by their form of words to see essences where there are none.

The very first sentence of a recent paper by Leshin, Leslie, and Rhodes (2020), tells us that "a problematic way to think about social categories is to essentialize them -- to treat particular differences between

people as marking fundamentally distinct social kinds” and having presented the results of a study that suggests that generic language promotes essentializing beliefs, concludes “avoiding generic language may be one way to limit the spread of essentialist beliefs”. It goes on to present evidence that even hearing innocuous generics about social kinds is thought to increase one’s chances of representing that kind as having an essence (whether it has one or not), and in turn invoking that essence in explanations of property possession and behavior of kind members (again whether such an explanation is plausible or not).¹¹

Most of this work has focused on the modal conclusion that the given kind has the given property by virtue of something inherent to the kind (essences, capacities, powers, etc). However, this again may only be one form of modal error. Given that the modal strength of generalizations can vary, amplification might occur in strengthening the generalisation itself -- for example, from something that only generalises over actual individuals to one that generalises to possible individuals in close possible worlds to one that generalises to all possible individuals. Although we haven’t presented -- and indeed won’t give -- an argument for this further conclusion, we can note that a priori it is reasonable, because the idea that generics involve, somehow, a modal element is quite well-represented in the literature. It seems we need to invoke goings on at other worlds to handle uninstantiated kinds, generics that have modal import but are not about kinds, or again sentences like

(9) Mail from Antarctica is sorted in this room

Which can be true even though no mail from Antarctica has ever arrived at this particular sorting center.¹² Moreover, a popular view of generics, defended most recently and forcefully by Nickel (2008, 2016) has it that the semantics of generics involve reference to normal individuals, where this is a modal notion. If Nickel is right, then given that we’re bad with generics, we’re probably bad with normality, and so prone to modal errors, including amplification.

Normative amplification.

A further species of amplification is normative amplification where subjects mistakenly amplify from acceptance of a non-normative generalisation to a normative one, or from a weak normative generalisation to a strong one. In addition to descriptive generalisations, generics often express or convey normative information. Moreover, generics vary in the types of normative information involved: they can articulate legal rules, moral norms, social norms, prescriptions, admonishments and contingent rules. Examples include:

- (10) Boys don’t cry.
- (11) Bishops move diagonally.
- (12) Children cross the road on the left.
- (13) Murder is wrong.

¹¹ See also Noyes & Keil (2020) for more nuances.

¹² For some discussion about modality and generics, and strategies for dealing with it, see Sterken (2015, 2017).

This means that generics potentially give rise to a variety of normative errors. Haslanger (2014) nicely outlines how this works for the more concerning cases like (10):

...in contexts where it is assumed that what's natural or good (at least for good things) is how things should be, that is, where such assumptions are part of the common ground, then the utterance of a generic enables a short inference to the normative conclusion, giving the generic a kind of normative force. If things ought to express their natures, and the utterance of statements such as 'Boys don't cry' implicates that, by nature, boys don't cry, then it follows that boys ought not to cry and we should reinforce boys' nature not to cry so they are not "deformed". (2014: 367)

This is both a plausible and important observation. It also finds some support in the empirical literature. Steven O. Roberts, Susan A. Gelman, and Arnold K. Ho (2016) report the results of an experiment which shows children have a tendency to infer prescriptions about how things should be from descriptions about how they are. The children were presented with novel kinds (e.g., Hibbles and Glerks) performing certain (morally neutral) actions and not performing others (say, being presented with Hibbles eating blueberries and avoiding strawberries).¹³ On being further shown behaviour not in accord with what they'd seen before -- for example Hibbles eating strawberries -- they tended to judge that as 'not okay': given some facts about how berries were eaten, they went on to conclude facts about which berries *should* be eaten. Generics, then, have a tendency to make us slide from *is* to *ought*, amongst other normative errors, and as the example above shows, such errors can have socially and politically problematic consequences.

Degradation

The next schematic form of error is degradation, which can be stated as follows:

Degradation: Subjects have a tendency to fail to accept a generic claim even when there is strong evidence for it, or accept the claim, but then later take the generic claim to have comparatively weak conclusions.

Though we're not aware of empirical work on this, besides amplification, there are plausibly statistical, modal and normative versions of degradation as well: in such cases, the acceptance of a relatively strong generalisation is warranted and appropriate, yet subjects nonetheless fail to accept any generalisation, or accept or conclude a relatively weak generalisation instead. Given that some of our thought is implicitly motivated and that acceptance of a strong generalisation may not suit our purposes, it seems plausible that sometimes subjects will deflate the strength of a generalisation when a stronger generalisation would be warranted, but it doesn't suit their purposes. As examples, consider:

(14) The privileged owe reparations.

(15) Lottery tickets are bad investments.

¹³ In the actual experiment, the berries were indicated by pointing, not naming, but this embellishment shouldn't matter for our point and makes describing it easier.

(16) Dogs are safer on a lead.

Imagine, for instance, how a privileged subject might represent the content of the generalisation in (14) as weaker than what they know there is evidence for. Likewise, there is overwhelming evidence for a strong interpretation of the generalisation in (15), however, if a subject is someone that loves buying lottery tickets or is generally a risk taker, then they may later think to themselves “yeah, some lottery tickets are bad investments, but not this one!” Many of us have a natural inclination to want to be right, are adverse to recognising our mistakes and don’t like correcting ourselves, so in circumstances where a well-supported generalisation conflicts with our deeply held beliefs, we’ll be more likely to infer something weak from it. For example, for (16), there might be overwhelming and undeniable evidence presented to me that dogs are safer on a lead, but I’ve been walking my dog without a leash for five years and I see how happy she is, so I later conclude from the generic that it is only a few rambunctious and silly dogs that are safer on a lead.

Once we become aware of the phenomenon, examples in the wild aren’t hard to spot. Here’s a disgusting one. Scientific American published an article¹⁴ ‘okay, so some dogs eat poop’. In it, they freely admit

(17) Dogs eat poop

Indeed, the evidence is pretty strong: a survey of 1,500 dog owners revealed that 62% of their pets ate poop daily, and 38% weekly. That, we think, clearly suffices for the truth of the generic. But notice that the headline is an existential quantification -- in the face of the evidence they themselves assembled, the writers weakened what they said. And they did so for obvious reasons: it’s gross.

We hardly think Scientific American are alone in that sort of rhetorical softening in the face of unpleasant facts, and so, although further evidence is needed, we think there’s a case for us to be on the look-out for degrading generics.

Indeed, we can perhaps bolster the case with this simple and a priori argument, that goes as so: amplification, both statistical and modal, clearly exists. But amplifying involves incorrectly inflating the strength of the generic. Consistency requires that if one accepts a given claim, then one rejects its negation. So consider the negation of an amplified generic:

(18) Women aren’t submissive.¹⁵

This is true, and the experience of most all of us will bear it out: there are many non-submissive women about. But if we hold the amplified positive claim, and we want to be consistent, we can’t hold ‘women aren’t submissive’, even in the face of many non-submissive women, and so we’ll be tempted to fall back on a weakened reading of it or perhaps even to an existential generalization. That is to say, we will be

¹⁴ <https://blogs.scientificamerican.com/dog-spies/okay-so-some-dogs-eat-poop/>

¹⁵ We can also use antonyms rather than negations: ‘women are independent/free/self-determining etc.’

tempted to degrade ‘women aren’t submissive’ and more generally whenever we amplify one generic we’ll need, on pain of inconsistency, to degrade its negation.¹⁶

Shifting

By shifting, we mean the following:

Subjects have a tendency to accept a generic claim based on certain evidence, but then later shift to accept another generic claim which is not supported by that evidence.

This probably sounds, and indeed probably is, pretty similar to the above mistakes of amplification and degradation, although as we’ve categorized things here, amplification and degradation are errors of inflating or deflating along some degree of strength scale -- quantity, modality or normative that corresponds to the strength of the operative generalisation, its underlying support or information about the kind. As we’ve presented them, one might think that if statistical amplification is going on, then modal degradation can’t be, for any combination of the errors, however, it could be that such errors can happen simultaneously or that we can have sequences of them. However, even simultaneously inflating and deflating along all these dimensions, perhaps along with other sorts of inference, may not capture all the shifting possibilities. Shifting allows subjects to shift the type of generalisation and doesn’t involve going up or down degree scales. Thus, a subject might in one context take the generic, ‘women are caring’ to support a generalisation like ‘most women have the capacity to be caring’, while in the next take it to support ‘all women should be caring in virtue of them being women’. Why wouldn’t this be captured by some permutation of inflating and deflating, plus perhaps some inference? First, perhaps it needn’t involve inference, and perhaps is just basic retrieval or interpretation in different contexts. Second, observe that the first generalisation ‘most women have the capacity to be caring’ is a capacity modal and is a descriptive generalisation, while ‘all women should be caring in virtue of them being women’ is deontic modality, is a normative generalisation and makes claims about the kind woman and the nature of that kind. While it might be possible to permute inflation and deflation, and make some inferences across many contexts to arrive at this normative generalisation, it’s far from obvious. Likewise, the communicative variant of this error seems plausible: that we often freely and unwittingly shift from interpretation to interpretation in different contexts. Thus, pure unmitigated shifting may be a genuine possibility.

Entrenchment

Entrenchment comes in two forms, which are mutually supportive. We state the two versions as follows:

Resistance: Subjects have a tendency to be *resistant* to giving up generic beliefs once accepted, despite evidence that they should be given up.

¹⁶ This is surely defeasible. The semantics and pragmatics of negated generics is a complicated topic we can’t get into here.

Stickiness: Generic beliefs themselves have a tendency to be *sticky*; they have a tendency to stick around through processes of belief revision where they would be candidates to be removed from the corresponding belief set.

We'll discuss these in turn.

Resistance

Why think that there is this resistance? Well, we can just note that, in fact, it happens. In the face of significant counter-examples, a subject will nonetheless not retract their generic belief. Take the simple example: *Birds fly*. Subjects retain belief in this generic despite the fact that penguins, ostriches and other birds don't fly. Generics are often observed to be difficult to refute or falsify. Subjects will engage in defensive shifting (Anderson, Haslanger and Langton 2012, Bosse 2021) where they shift from interpretation to interpretation in order not to avoid having to give up their generic belief. Anderson, Haslanger and Langton (2012) provide an example:

Attending to this slippage between majority and characteristic readings helps us explain certain shifts in the defense of generics. Consider [...]:

[19] Latinos are lazy.

Does [19] assert a majority generic or a characteristic generic? Interpret [19] as a majority generic. To combat it, one provides many counterexamples. However, the speaker can then suggest that, although many Latinos aren't lazy, they tend to be, thus embracing the characteristic generic. Instead interpret [19] as a characteristic generic. To combat it one provides evidence that, say, Latinos show no greater tendency towards laziness than any other group. The speaker can then suggest that, although it is not part of the nature or essence of Latinos to be lazy, most are. This slide back and forth between different interpretations of the utterance allows speakers to avoid taking responsibility for the implications of their claims. (764, numbers changed)

Such shifting also indicates that subjects don't even seem to know what would count as decisive counter-evidence to their generic beliefs.

Such examples support *Resistance*, as some of the blame for the stickiness here must lie with the subject themselves as it is sometimes willful or intentional. But they also support *Stickiness* as it is plausibly representational properties of generics that bear some of the responsibility (e.g., the shifting might be semantically sanctioned).

Stickiness

Some generic beliefs stick around because they're stably part of the common ground (for which see e.g. Stalnaker 2002, among much else), we're not attending to the fact that they're there (Haslanger 2011), and they're so widely accepted that giving them up is either not under our control or would cause us

difficulty in coordinating with other people and in avoiding confusion because we don't share common ground.

Haslanger (2011) and Langton (n.d.) draw attention to cases where some of the content conveyed by generics is added to common ground without our recognition of it. Haslanger's version of the argument goes as so. As already noted, she thinks, plausibly enough, that generics implicate claims about essences which, in light of reality (there aren't the relevant social essences) are false. But, notoriously, implications (as well as other not-at-issue content), even false ones, can get into the common ground. This is because the common ground is something constructed and updated in order to make sense of what the speaker is saying. If the only way to make sense of what a speaker is saying is by adding to the common ground some not-at-issue content about social essences, we might do this accidentally, in roughly the same way that we find ourselves accepting what we'd rather not accept when faced with sneaky questions like 'Have you stopped stealing from children's charities?', any answer to which will shift the common ground in ways we don't want (by adding to it the proposition that I formerly stole from children's charities). By exploiting the essentializing presuppositions or implicatures of generics, then, speakers can cause harmful claims to become part of the common ground in a way that is difficult to avoid.

In the same vein, some generic beliefs are so widely shared that they play a vital role in our ability to coordinate socially, epistemically and communicatively, and avoid confusion. Take, for instance, *guns are dangerous*, which has all the properties just listed, because of which it would be difficult to give up.¹⁷ One of the reasons generic beliefs behave this way lies in the nature of generic representation itself -- the flexibility and the placeholder-like, semantically underdetermined or non-specific qualities of the generic quantificational operator *Gen*, allows subjects to hold on to a placeholder, semantically underdetermined or non-specific generic belief which can be used to infer a large variety of differing more specific generalisations in context (witness the examples above). Basically, a crucial capacity of generic representation is that subjects can share this generic place-holding, underdetermined, non-specific belief while they might not share the more specific occurrent ones. Structurally this can occur in terms of shared belief across a population of subjects, but also across the stages of a single subject, so that the beliefs are socially sticky, in addition to being cognitively sticky. Empirical work by Cimpian & Scott (2012) suggests that subjects assume that generic knowledge is widely shared, so it is not simply the representational properties of generics but also some form of higher order confirmation that makes generic beliefs sticky. *Resistance* and *Stickiness* also reinforce one another, since subjects have a tendency to be resistant to giving up generic beliefs, as evidenced by the phenomena of being hard to refute and to disagree with, which make it hard to challenge the beliefs that are sticking around. Generics give an illusion of agreement, both personally and interpersonally, which is potentially detrimental to productive epistemic and social progress. All this is just to say: generic beliefs have a tendency to become cognitively and socially entrenched, and in certain circumstances this will be beneficial, but in others detrimental.

¹⁷ *Guns are dangerous* is good for children to know, for governments to know, for movie-makers to know, for soldiers to know, for gun-club owners to know, but they each will understand it in a different way and act on that understanding in different ways, but even despite its different understandings, we want different sectors of society to coordinate on the generic.

1.3 Generic assertions can cause or constitute harm

As discussed in the introduction generics, like (3) and (4), can cause or constitute harm when asserted.

When Trump uttered (5), it wasn't only offensive and reprehensible because it was false and politically incorrect. It was offensive and reprehensible because its content and delivery was tantamount to slurring. Likewise, for utterances of (3) and (4). On many views of slurs (e.g. that of Hom, in the paper already cited), the source of their offensiveness is that the slur expresses or communicates negative stereotypes of the slur's target group. Generics are the primary vehicle of stereotypes (and indeed the vehicle most invoked in stereotype semantics of slurs) and hence, when such generics are explicitly uttered, they possess some of the effects of slurs. Uses of slurs cause or constitute harm by demeaning and dehumanizing members of the slur's target group (Jeshion 2013). Likewise, uses of such generics cause or constitute harm by demeaning or dehumanizing members of the given kind. Such uses of generics, indisputably should be avoided. But it's unclear if avoidance of their use alone justifies any form of general avoidance.

However, if stereotyping often causes harm and generics are the primary vehicle of stereotypes (Hammond and Cimpian 2017), then this might be reason enough to avoid them – in order to limit stereotyping. Or, if uttering generics is implicated in the reinforcement of problematic ideologies (as in Haslanger 2011) so that their lexical effects include a contributory role in the harms of problematic ideologies, then this might be reason enough to avoid them.

The harms caused by stereotyping or supporting problematic ideology needn't harm individual targets directly in their context of utterance or thought. Uttering stereotypes, even seemingly quite innocuous ones, can serve to reinforce and promote socially and politically problematic structures and ideologies (which, in turn, may harm their targets). Take again utterances of (6) from the section on modal amplification as an example. Hearing (6) promotes and reinforces ideologies involving gender. Even seemingly beneficial slogans like *girls can do whatever boys can* might reinforce the idea that boys set the standards of what it is appropriate or good to do (Chestnut and Markman 2018). Even negated generics like *girls don't spend lots of time at the mall* support inferences to the effect that there is something inherent to the kind girls that explains their properties and behaviors (Foster-Hanson, Leslie, and Rhodes 2016).

1.4 Generics make speakers especially prone to communicative and interactional issues

Another way in which generics might be hazardous is that they make speakers especially prone to misinterpretation and miscommunication. Because much of the content of generics is implicit, much is left open for the interpreter and it is, then, difficult for the speaker and hearers to successfully share

content. We discussed many cognitive errors above, and these all have correlates at the communicative level. A speaker might utter a generic and intend for her hearer to interpret it as a relatively weak generalisation, but instead the speaker misinterprets the generalisation expressed as something with a strong quantificational, modal or normative force. We won't go into all the variations on the ways in which speakers might fail to share content or coordinate in communicative contexts. We'll merely note the propensity caused by generics, for even well intentioned speakers to be misled into thinking that they've successfully shared content when in fact they haven't.

However, this is just to note the potential hazards of using generics amongst well-intentioned speakers. The hazards are much worse upon consideration of ill-intentioned speakers. Ill-intentioned speakers can exploit the implicit, underdetermined nature of *Gen* by hiding what was really intended, by masking changes in their views and maintaining plausible deniability (cf. Saul 2017), amongst other things.

Now that we have a good sense of the many ways in which generics might lead us astray semantically, cognitively, epistemologically and communicatively, and some reasonable motivation for avoiding them, let's consider why not all generics can or should be avoided.

Section 2: Why not all generics can or should be avoided

We've just seen many solid reasons for worrying about generics. Taken together, we agree with the authors such as Anderson, Haslanger, Langton (2012), and Leslie (2017), as well as the many psychologists whose important work we cited that there is, at least, a strong *prima facie* case for avoiding generics, in particular those involving racial and gender kinds.

However, that case is not beyond question, and indeed some authors have questioned it. In this section, we will present previous work presenting difficulties for any sort of general prescription for the avoidance of generics. In the one following it, we will consider some further candidate prescriptions that so far have not been considered.

We think we can find, either implicitly or explicitly, at least four reasons for being doubtful about any sort of prohibition of generics: they are what we'll call *overgeneration*, according to which the prohibition overgenerates, *hermeneutic injustice*, according to which without generics we'd be unable to speak about and thus make sense of important features of social reality, *impossibility*, according to which expunging generics from our language can't be done, and *backfiring*, according to which banning generics would have unintended bad consequences such as the increased use of code words and other language that can be used to sneakily marginalize people.¹⁸

Overgeneration

¹⁸ See Khoo, this volume, and Saul, this volume for examples.

Jennifer Saul (2016) makes the following very important observation that can be used to cast doubt on prohibition: when talking about topics like gender and race, generics aren't necessary to fall victim to or enable some of the cognitive mistakes we pointed to above.

In particular, think of what we called modal amplification, one important realisation of which was the tendency to essentialize: on the basis of a generic generalization to assume that it's of the nature or essence of the kind generalized about to be that way. Saul points out we don't need generics to modally amplify.

She tells this story:

A local shopkeeper, who spotted my “This is what a feminist looks like” bag... decided to tell me some of her supposedly feminist beliefs. She proceeded to utter [20] and [21].

[20] Many Asian men are abusive to women.

[21] 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... there would seem to be the problematic implicature that being Asian is relevant to the dangerous property being attributed. (Saul 2017:6, numbering changed)

But, as Saul goes on to note, (20) and (21) do not involve generics, but seem clearly to carry an essentializing message according to which Asian men are bad to their wives. There's accordingly a very natural overgeneration worry: if we have a case for prohibiting generics, then we have a case for prohibiting (other)¹⁹ quantifiers.

We can bolster the case. Experimental work that Saul, along with Hoicka, Prouten, Whitehead and one of the authors of this paper (Sterken) carried out seems, among other things, to bear out Saul's observation. That work showed that children and adults presented with sentences containing quantified noun phrases, where the quantifier was high-proportion, and the noun phrase novel ('Zarpie') encouraged essentializing, suggesting that the propensity to essentialize can't be tied exclusively to the generic, and implies that the abolition of generics won't abolish the propensity to essentialise.

Hermeneutic injustice

Hermeneutic injustice is the by now familiar phenomena brought to prominence by Miranda Fricker according to which certain people are disadvantaged because they lack the conceptual resources necessary to make sense of their experience.

¹⁹ Ignore the parenthesis if you don't think generics are quantified.

We think that work by Katharine Ritchie can be understood to imply that a prohibition on generics would lead to hermeneutic injustice. Let's begin by just quoting her:

Racial, gender, and other social generics can be useful because they accurately describe systematic patterns of injustice. An apt description of structural oppression requires capturing that it is widespread, general, and systematic. Generics capture general structural patterns in a way that overtly quantified statements do not. (2019: 38)

There are a few claims here. First, that there are systematic and structural patterns of oppression in society, and perhaps that these are especially prevalent when it comes to race and gender. The second is that the truth²⁰ of generics “rely on a group having an intrinsic, persistent, and causally explanatory nature”. The third, implicit and more questionable claim, is that only generics rely on this.

To spell this out, consider:

(22) Black people face economic, legal, and social discrimination

We should accept that this gives voice to a pattern of structural oppression. And given the above, we should also accept that the generic somehow conveys that the denotation of the subject term satisfies the predicate in some sort of non-accidental way. Arguably, these two facts go together pretty well: the latter explains why we understand the former as we do.²¹ Generics are that bit of language that enable us to talk about structural oppression, and so since we want to talk about structural oppression, we should retain them.

There's definitely something to be said for this. Even if, as the above section suggests, generics aren't uniquely apt for essentializing, they are, we don't deny, a tool for doing so, and so we think Ritchie's argument deserves serious consideration, even if, because of the falsity of the implicit claim that *only* generics are apt for essentializing, we don't think it is decisive.

Impossibility

By *impossibility*, we mean that it's unlikely, even if we were convinced that removing generics from communication was a good idea, that it's an idea we could implement.

Part of the source of this is simply the general worry, familiar from the literature on conceptual engineering for example, about how it's hard to change language. Thus for example we might wish that ‘woman’ meant something that highlighted women's subjugated social role, but most English speakers

²⁰ We don't need to commit ourselves to any claims about truth. The point made goes through if we replace ‘truth’ with ‘typical understanding’ or something like that.

²¹ Ritchie claims that the corresponding ‘many’ claim fails to convey the sort of essentialization needed to make sense of what the sentence says. We disagree somewhat with her here: at least we don't think this example can be generalized from.

aren't attentive to philosophers' prescriptions (nor are influencers who arguably *do* have some slight ability to affect speakers' languages. If Taylor Swift suggested everyone use the Haslangerian ameliorated definition of 'woman', some Swifties would probably give it a go. But, as far as we know, Swift isn't a Haslanger-reader.)²²

That, we think, is fair enough. But it's a domain-neutral point: it applies to any attempts to change language. It's worth inquiring whether there are reasons to think generics in particular will be resistant to attempts to ameliorate them.

In order to do so, it will be useful to consider again the authors we've discussed who have put forward proposals to change our generic talk. While reminding the reader that our aim in this paper is not to criticise any one particular author's take, but rather to explore the issues themselves, we can note that for Leshin, Leslie, and Rhodes:

avoiding generic language may be one way to limit the spread of essentialist beliefs and avoid some of their negative consequences for inter-group relations (2020: 26)

By contrast, Anderson, Haslanger and Langton (2012) were much concerned with changing the way we respond to others' use of (racial) generic talk.

This can help us draw a pair of distinctions: we can distinguish *removing generic language itself*, and *changing how we respond to it*. Prima facie, there's no reason to think there's any dependency between these two: maybe it's impossible to remove generics but possible to change how we react to it (admittedly, we can't remove generics while keeping how we react to them the same).

So, a slightly narrower question is: should we think it's possible either to get rid of generic language itself or how we react to it?

We think both are doubtful, even abstracting from the more general considerations about changing language we adverted to at the start of this subsection. Our reasons aren't hugely interesting. For one, although you might initially be more hopeful that your responses to uses of language are something more under your control than the public language itself, this line is specious. For any large scale change to happen, it would need to be that the responses of many, perhaps most, language users change, but again we think this is no more possible than changing people's use of 'woman'. Any large scale coordinated action is difficult, and responding is no different.

²² Actually, this point requires a bit of care. One of the contributors to this volume, Dembroff, was an amicus curiae to a recent Supreme Court case about discrimination on the basis of attraction and gender, called upon because of their expertise in the philosophy of gender and sexuality. While the Supreme Court is not Taylor Swift (in many respects), it is an important institution which plausibly has a lot of influence when it comes to changing meanings. For some information about this case, see:

<https://news.yale.edu/2020/07/08/yale-philosopher-and-legal-scholar-collaborated-recent-scotus-case>

Moreover -- and this will seem pedantic, but it will help us get one of our positive proposals on the table -- we don't fully see what such exchanges would look like. Anderson, Haslanger, and Langton ask us to dissent from even true statistical generics because they implicate an (always false) characterizing generic. But that will make conservation very difficult! Consider this stultifying exchange, between an American A in Dublin for the first time and a native interlocutor.

A: Irish people sure do like poor quality grocery store Chicken sandwiches.

The interlocutor, B, can't say, 'No, Irish people don't like those sandwiches', because that contains a generic. So it seems they'll be compelled to say something like:

B: That's false. / You're wrong. / No.

A reasonable response for A would then be:

A: Oh, really, so they don't like them?

Assuming that 'they' inherits its genericity from its antecedent, again it seems B must deny, saying

B: That's false.

We're tempted to say that A will come away from this dialogue, not with an appreciation for the perniciousness of generics, but with a negative impression of B's conversational abilities and/or personality, and so we don't think the Anderson, Haslanger, and Langton line is yet clear. That said, we like and will take inspiration from the idea that we should try to disrupt generic talk, when we later discuss transformative communicative disruption in section 3.

There are also reasons for thinking that generics in particular will be hard to eliminate from our language. In fact, we think the very same empirical research that people like Leshin, Leslie, and Rhodes appeal to makes this case. If generics are indeed cognitively fundamental, more easily and earlier grasped and remembered than other devices of generalization, then getting rid of them, even on an individual level, seems like it will be an uphill battle, akin to getting rid of our propensity for sugar. We can try, and we might go a couple of days without generics, but they're just too cognitively fundamental to extirpate from our own idiolects, to say nothing of the natural languages in which they occur.

Backfiring

Let us begin this final possible reason against a generic prohibition with a(n ugly) story. An alt-right podcast which we won't name would add an echo effect when they mentioned a Jewish person's name. That got translated into text as a set of brackets, and

During the summer of 2016, Jewish journalists began noticing that alt-righters were typing their name in parenthesis -- (((like this))) -- in a written simulation of the sound effect. Some tagged

with such punctuation got an extra helping of anti-Semitism in the form of death threats and Photoshopped pictures of themselves in concentration camps. (Wendling 2018:83)

There are several points to note here. First, this indirectly supports Hoicka, Prouten, Whitehead, Saul and Sterken, who claim that essentializing isn't limited to generics. The alt-right made up a device to essentialize proper names in a negative way: in saying *((Michael Bloomberg)) supports multiculturalism*, one is essentially saying *Michael Bloomberg qua Jewish person supports multiculturalism*, with a heavy implication that it's his Jewishness that leads to his support. More generally, it makes us pessimistic that excising a bit of language would stop essentializing. People would move to other devices for doing so, and indeed might move to more covert, harder immediately to spot and root-out devices.

Of course, this council of despair only goes so far. After all, one could imagine the above line spoken to someone a couple of decades ago telling them to give up their campaign to remove the r-word from polite use. More generally, the advice that changing language is pointless because bigots will always adapt is, we think, bad advice, and we don't whole-heartedly endorse it. Rather, we should be alive to the chance that our efforts will backfire, and will do so in probably completely unpredictable ways (or at least ways that philosophers of language, as opposed maybe to sociolinguists, can say something interesting about).

Section 3: Must do better

To briefly recap: we've seen that generics are a locus of communicative and epistemic harms, which has led some theorists to think we should avoid using them, or respond to uses of them in ways that blunt their harm. We've also seen some reasons for doubting this, that mainly centered on the fact that generics aren't uniquely harmful and that anyway avoiding them is impractical or might have bad downstream effects. But a different, more positive, perspective is offered by Saul (2016), who argues that we should retain generics and *just get better* at using them. She writes:

... 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. (2016: 14)

Thus, Saul suggests we shouldn't stop asserting generics, rather we need to get better at reflecting on whether the generic generalisations we believe are warranted and whether the ones we encounter in the wild are worthy of believing. We also need to be good interlocutors and epistemic citizens and (i) only assert generalisations when we've suitably reflected on the evidence for it, and (ii) respond to the generic generalisations we hear with appropriate suspicion, as Saul (2017), Anderson, Haslanger and Langton (2012) and Haslanger (2011) suggest (in slightly different ways).

Saul's suggestion is a good one. However, it would be nice to have more guidance not just with respect to what sort of evidence gives warrant to a given generic generalisation (recall this is thorny and complex,

and we aren't yet particularly good at it), but also with how one is to know when it is ok to assert a generic generalisation given our evidence (especially when we now know our hearer is likely to misinterpret it, or that it may be stored long term in our belief set as a place-holder generic belief from which we can infer many things). We need to know more about the sorts of contexts we're in, including what kind of evidence we have, and what kinds of generalisations are relevant and warranted given the kind of context we're in. This will allow us to become better at thinking about, interpreting and asserting generics.

A general strategy, which we find worthy of more consideration, is to think more creatively about changing our communicative norms and practices. In particular, one suggestion, closely related to Saul's, would be to work to heighten the standards of precision in contexts where gender and racial kinds are discussed and generics asserted. (To understand what is meant by this: think about the standards of precision for generalisations in mathematical contexts vs. those in ordinary "shooting the breeze" type contexts.) This change would result in what Saul suggests, it would force speakers to reflect more on the evidence for their own assertions and those of others, and would make critique easier. It would also likely cause the form and content of generalisations to be more explicitly articulated; perhaps even over time the class of generalisations which are operative in our epistemic and linguistic community would become less blunt an epistemic resource (one can only dream). The trouble with this suggestion is that some of the issues discussed in Section 2 will come back to haunt us. We often want to encourage discourse about social kinds (e.g., to talk about race and gender with our children, to recognize the bad things so that they can be rectified, etc), and we want to encourage beneficial uses (e.g., à la Ritchie). But if we need to do the extra work of reflecting each time we hear or assert a generic or the standard of precision in those contexts are high, then we are putting a barrier in the place which discourages such discourse.

These strategies, though initially promising, may not be the best we can do. One outcome of the burgeoning literature on conceptual engineering is recognition of the wide variety of strategies one might take to implement representational changes, similar in kind to the ones we've been discussing. Instead of eliminating generics as the prohibitionist might have it, let's stick with the spirit of more conservative strategies just discussed and attempt to improve the situation with generics. The literature on conceptual engineering is rife with examples upon which we can take inspiration. We'll briefly consider some proposals for how to ameliorate other representational devices, and look for some analogies in connection with generics. We'll make three preliminary proposals along these lines, however, we'll have to save substantial discussion for another venue.

Prescription 1: Ameliorating Gen

Kevin Scharp (2013) has a proposal for how to ameliorate the concept of truth. He argues that the concept is defective because its constitutive principles are inconsistent. In order to eliminate these defects, he proposes that our defective concept of truth needs to be replaced by two consistent successor concepts *ascending truth* and *descending truth* that are to be used in different circumstances depending on the kind of theoretical or expressive work to be done.

Mimicking the structure of Scharp's ameliorative strategy, we might replace *Gen* with successor representations $Gen_1, Gen_2, \dots, Gen_n$ that serve the useful purposes to which we put *Gen*, but does so in a way that avoids as many of the problematic issues with *Gen* discussed in section 1. In formulating such a strategy in detail (something we cannot do here), we will be forced to formulate and decide precisely the guidance needed to get better at thinking and talking about generics. That is, we will need to specify in what circumstances characterising generic generalisations are needed and warranted; when it is that essentialising inferences are needed and warranted; and so on.

Is it feasible? Well, as feasible as other proposals. Although above, under *impossibility*, we gave some reasons for thinking that removing generics from our thought and talk would be difficult, we didn't say it would always definitely be impossible. And below we'll briefly allude to some cases in which it does seem that we've managed to shape language to our purposes. But even if the above proposal isn't feasible, it could still serve a useful purpose, even if just for us as theorists of generics and the pragmatics of language. Much productive work in 20th century philosophy was a result of idealizing, of thinking what a language should be like (for the purposes of science, for example). That work yielded things like Russell's theory of descriptions, and theories of scope and quantifiers and operators (for just a handful of examples) that have proved to be very useful lenses through which to understand natural language. Disentangling possible meanings *Gen* could have, then, could be useful even if in English as it stands it doesn't have those meanings.

Prescription 2: Special proposal for striking property/slurring cases and problematic entrenched cases

The above (and below) prescriptions attempt to make the most of generic language by improving it or foregrounding its positive features. But it is helpful, as a matter of strategy, to think about how to deal with particularly entrenched and problematic generics, especially as and when we encounter them in the wild today. Thus, consider again one of the particularly tricky sets of cases, that of striking property social generics, such as:

- (3) Muslims are terrorists

We think that people will continue to say and think such things, either explicitly to subjugate people, or less malevolently because of the psychological and semantic confusions generics breed. And we're not hugely confident what would succeed in such cases, so we think we need to find a way to deal with them as unavoidable but unfortunate features of social life.

One suggestion we like is hinted at in Anderson, Haslanger, and Langton and developed in more detail in Sterken (2020). In the former work, as we've seen, the authors propose the somewhat surprising strategy of rejecting even true generics because of false implicatures they have. Doing so, of course, violates quite fundamental rules of conversation. In particular (and assuming the generic sentence in question was relevant in the conversational context in question), it can be seen to do something like flout Gricean quality: it treats something true the way something false should be treated.

We suggested that this would lead to conversational impasses in ways that Anderson, Haslanger, and Langton didn't seem to appreciate. But a paper by one of us treats these impasses as a helpful tool for ameliorative language improvement.

The basic thought here is as follows. In the conceptual engineering literature, there is quite a lot of discussion about the (im)possibility of fixing language -- of actually making the changes we as theorists think it's desirable to make: people discuss whether such changes are possible, whether externalism in metasemantics renders it impossible, and so on. But another approach is to point out that, as a matter of fact, language *does* change, and see what we can learn from this fact. Language changes when we baptise someone newly born (say, a baby). It changes when we introduce new slang on social media, when we reappropriate slurs, and in other cases (Sterken 2020: section 2).

Now imagine a case when a term has just been introduced, but perhaps isn't completely widespread. One way to think of this is that in such cases, we can have no guarantee that when we enter a conversation, each of the participants' idiolects is synched up to the shared public language they speak.

Important questions about conversational and conceptual ethics arise in such cases. The different idiolects might lead to ostensible violations of maxims of truth or relevance or politeness. As an example of (roughly) the latter: a gay person from, say, New York, might offend a gay person from a less progressive place if they use 'queer' when only in the former place has the term been reappropriated.

Sterken (2020) argues that these mismatches can be used to improve language. As she writes:

Ameliorators, given the right motivations, can use deviant communication – like miscommunication, uncooperative or insincere speech – to accomplish their projects of linguistic change. The disruption of standard communicative patterns can help them accomplish their goals. The disruptions are a good thing as they can have the effect of making the hearer stop and reflect on their usage, and this reflection can be transformative. In other words, the deviant communicative activity of the ameliorator can engage the hearer in the sort of metalinguistic reflection needed to acquire the new meaning and understand the ameliorator's utterances as she intends. (2020: 430)

Now, the evidence marshalled so far in this paper should make it plausible that generics are sources of much deviant communicative activity. And so the Sterken perspective would enjoin us to make use of it to improve generics.

Here's a concrete case. Few, but not no, Muslims are terrorists. The racist uses this fact to be able to utter plausibly sentences like (2). The Anderson, Haslanger, and Langton line would tell us to reject such utterances. But an alternative is to try to disrupt the conversation. In response, one could say:

(23) Norwegians are terrorists.

We could imagine the conversation continuing:

A: That's ridiculous

B: What? Anders Breivik was responsible for the terrorist attack at Utoya. He's Norwegian.

A: That's just one example.

B: But it only takes a very small percentage of a group G to do something for you to say that Gs do that thing, right? After all, a very small percentage of Muslims are terrorists. You know that, right?

At this point, it seems to us, the conversation, and so the racist's aims, have been thoroughly disrupted. Hopefully, the racist will 'stop and reflect on their usage', and at the least be more careful in future.

We think this generalizes: if generics are a tool that racists can misuse, then anti-racists, rather than avoiding the tool, can simply use it in furtherance of their anti-racist goals. By uttering the falsity that 'Norwegians are terrorists' one might get the racist to retract their falsity that Muslims are terrorists. One way to treat the misuse of generics, then, is to fight fire with fire.

Prescription 3: The "beneficial generics" as linguistic practice strategy.

It's sometimes observed, as it is above, that generics have beneficial uses -- encouraging, recognitional or counterspeech cases such as *girls are tough* or (22) from above. We might take inspiration from these cases and take them as a basis for a general strategy of communication with generics. Why not try to use beneficial generics as much as possible, not just in contexts of encouragement, recognition and counterspeech? Such cases act as a corrective of people's lagging social beliefs and of the social world. If we used generics as corrective across the board, then perhaps we'd not only address some of the concerns above, but our linguistic practices with generics would be proactive. They would address problems with people's social beliefs and help promote and reinforce a better social reality.

On this strategy, speakers would go around saying things like *girls are tough*, *black people are good at math*, *bankers are kind and compassionate*, etc. Basically, our linguistic practices involving generics, not mere counterspeech or one off beneficial uses, would play a proactive role as corrective of social reality and poor ideology. Utterances of *Girls are tough* and *black people are good at math* give recognition where it is due, and play a role in correcting lagging social ideology, which can go some way to remedying our social situation. We might even utter generic generalisations that are not yet true, but we hope might become true. Perhaps if we all went around saying that *bankers are kind and compassionate*, then maybe that's what bankers would be.

This would address some of the concerns above: errors such as essentializing and statistical amplification could be working for us, rather than against us. However, this strategy, for the familiar reasons above, could not be taken across the board. It wouldn't work for cases like *immigrants have access to great schools*, where recognition of the problem is crucial to making social progress. It may also backfire in the sense outlined in section 2. However, perhaps some restricted version of the strategy can be found. If the

strategy in Prescription 1 works, then we might learn more about the contexts in which it would work, and can restrict this strategy accordingly.

References

- Abelson, R. P., & Kanouse, D. E. (1966). Subjective acceptance of verbal generalizations. In *Cognitive Consistency* (pp. 171-197). Academic Press.
- Anderson, Luvell ; Haslanger, Sally & Langton, Rae (2012). Language and Race. In Gillian Russell & Delia Graff Fara (eds.), *The Routledge Companion to Philosophy of Language*. Routledge.
- Anderson, Luvell & Lepore, Ernie (2013). Slurring Words. *Noûs* 47 (1):25-48.
- Anderson, Luvell (2017). Hermeneutical Impasses. *Philosophical Topics* 45 (2):1-19.
- Baldwin, D. A., Markman, E. M., & Melartin, R. L. (1993). Infants' ability to draw inferences about nonobvious object properties: Evidence from exploratory play. *Child development*, 64(3), 711-728.
- Barwise, Jon, and Robin Cooper (1981). "Generalized quantifiers and natural language." In *Philosophy, language, and artificial intelligence*, pp. 241-301. Springer, Dordrecht.
- Bosse, Annie . In progress. Generics in use. Manuscript, University of Cambridge.
- Cappelen, Herman (2018). *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.
- Cappelen, Herman & Dever, Josh (2019). *Bad Language*. Oxford University Press.
- Carlson, G. N. (1977). Reference to kinds in English. University of Massachusetts, Amherst.,
- Carlson, Greg N. & Pelletier, Francis Jeffrey (eds.) (1995). *The Generic Book*. University of Chicago Press.
- Chestnut, E. K., & Markman, E. M. (2018). "Girls are as good as boys at math" implies that boys are probably better: A study of expressions of gender equality. *Cognitive science*, 42(7), 2229-2249.
- Chierchia, Gennaro (1998). Reference to Kinds across Language. *Natural Language Semantics* 6 (4):339-405.
- Cimpian, A., & Scott, R. M. (2012). Children expect generic knowledge to be widely shared. *Cognition*, 123(3), 419-433.

Cimpian, Andrei ; Brandone, Amanda C. & Gelman, Susan A. (2010). Generic Statements Require Little Evidence for Acceptance but Have Powerful Implications. *Cognitive Science* 34 (8):1452-1482.

Cimpian, A., Mu, Y., & Erickson, L. C. (2012). Who is good at this game? Linking an activity to a social category undermines children's achievement. *Psychological science*, 23(5), 533-541.

Foster-Hanson, E., Leslie, S. J., & Rhodes, M. (2016). How Does Generic Language Elicit Essentialist Beliefs?. *CogSci*.

Gelman, S. A., Leslie, S. J., Was, A. M., & Koch, C. M. (2015). Children's interpretations of general quantifiers, specific quantifiers and generics. *Language, cognition and neuroscience*, 30(4), 448-461.

Gelman, S. A., Tapia, I. S., & Leslie, S. J. (2016). Memory for generic and quantified sentences in Spanish-speaking children and adults. *Journal of Child Language*, 43(6), 1231-1244.

Gelman, S. A., Ware, E. A., Manczak, E. M., & Graham, S. A. (2013). Children's sensitivity to the knowledge expressed in pedagogical and nonpedagogical contexts. *Developmental psychology*, 49(3), 491.

Graham, S. A., Kilbreath, C. S., & Welder, A. N. (2001). Words and shape similarity guide 13-month-olds' inferences about nonobvious object properties. In *Proceedings of the twenty-third annual conference of the cognitive science society* (pp. 352-357).

Hammond, M. D., & Cimpian, A. (2017). Investigating the cognitive structure of stereotypes: Generic beliefs about groups predict social judgments better than statistical beliefs. *Journal of Experimental Psychology: General*, 146(5), 607.

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

Hoicka, Elena, Jennifer Saul, Eloise Prouten and Laura Whitehead (nd). Language Signalling High Proportions, not Just Generics, Leads to Essentializing for Novel Social Kinds. Preprint here: <http://rachelsterken.org/quantifiers%20vs%20generics%20jcep%202020%2005%2013.docx>

Hom, Christopher (2008). The Semantics of Racial Epithets. *Journal of Philosophy* 105 (8):416-440.

Jeshion, Robin (2013). Expressivism and the offensiveness of slurs. *Philosophical Perspectives* 27 (1):231-259.

Krifka, Manfred (2003). Bare NPs: Kind-referring, Indefinites, Both, or Neither? *Semantics and Linguistic Theory* 13:180.

Lakoff, G. (2004). Don't think of an elephant: Progressive values and the framing wars—a progressive guide to action. White River Junction, VT: Chelsea Green Publishing.

- Langton, Rae (1993). Speech acts and unspeakable acts. *Philosophy and Public Affairs* 22 (4):293-330.
- Langton, Rae & West, Caroline (1999). Scorekeeping in a pornographic language game. *Australasian Journal of Philosophy* 77 (3):303 – 319.
- Langton, Rae (nd). iv. Silence As Accommodation Failure, lecture notes, available at:
http://media.philosophy.ox.ac.uk/assets/pdf_file/0005/37319/HANDOUT_Lecture_IV._Silence_as_Accommodation_Failure.pdf
- Leslie, Sarah-Jane (2008). Generics: Cognition and acquisition. *Philosophical Review* 117 (1):1-47.
- Leslie, Sarah-Jane (2017). The Original Sin of Cognition: Fear, Prejudice, and Generalization. *Journal of Philosophy* 114 (8):393-421.
- Leslie, S.J., & Lerner, A. (2016). Generic Generalizations. In E. Zalta (ed.), *Stanford Encyclopedia of Philosophy*
- Leslie, S. J., & Gelman, S. A. (2012). Quantified statements are recalled as generics: Evidence from preschool children and adults. *Cognitive psychology*, 64(3), 186-214.
- Leslie, S. J., Khemlani, S., & Glucksberg, S. (2011). Do all ducks lay eggs? The generic overgeneralization effect. *Journal of Memory and Language*, 65(1), 15-31.
- Leshin, R., Leslie, S. J., & Rhodes, M. (2020). Does it matter how we speak about social kinds? A large, pre-registered, online experimental study of how language shapes the development of essentialist beliefs.
- Liebman, David (2011). Simple Generics. *Noûs* 45 (3):409-442.
- Prior, A. N. (1960). The Runabout Inference-Ticket. *Analysis* 21 (2):38.
- Nickel, Bernhard (2008). Generics and the ways of normality. *Linguistics and Philosophy* 31 (6):629-648.
- Nickel, Bernhard. 2010. "Generically Free Choice." *Linguistics and Philosophy* 33 (6) (December): 479–512. doi:10.1007/s10988-011-9087-4
- Nickel, Bernhard (2016). *Between Logic and the World: An Integrated Theory of Generics*. Oxford University Press UK.
- Noyes, A., Keil, F. C. (2020). There is no privileged link between kinds and essences early in development. *PNAS*.

Pelletier, F.J. (2010). A Philosophical Introduction to Generics. in F.J. Pelletier (ed) *Kinds, Things and Stuff* (OUP): 3-15.

Ritchie, Katherine (2019). Should We Use Racial and Gender Generics? *Thought: A Journal of Philosophy* 8 (1):33-41.

Ritchie, Katherine & Knobe, Joshua (forthcoming). Kindhood and Essentialism: Evidence from Language. In Marjorie Rhodes (ed.), *Advances in Child Development and Behavior*.

Roberts, Steven O., Susan A. Gelman, and Arnold K. Ho. "So it is, so it shall be: Group regularities license children's prescriptive judgments." *Cognitive Science* 41 (2017): 576-600.

Rothbart, Myron, Solomon Fulero, Christine Jensen, John Howard, and Pamela Birrell. "From individual to group impressions: Availability heuristics in stereotype formation." *Journal of Experimental Social Psychology* 14, no. 3 (1978): 237-255.

Saul, Jennifer (2017). "Are generics especially pernicious?." *Inquiry*: 1-18.

Stalnaker, Robert (2002). Common ground. *Linguistics and Philosophy* 25 (5-6):701-721.

Stanley, Jason (2015). *How propaganda works*. Princeton University Press.

Sterken, Rachel Katharine (2015). Generics in context. *Philosophers' Imprint* 15.

Sterken, Rachel Katharine (2017). The Meaning of generics. *Philosophy Compass* 12 (8):e12431.

Westerståhl, Dag (2019). Generalized Quantifiers. *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.), URL
=<<https://plato.stanford.edu/archives/win2019/entries/generalized-quantifiers/>>.

