Giving Generic Language Another Thought[1,2]

Eleonore Neufeld, Anne Bosse, Guillermo Del Pinal, Rachel Sterken

Abstract

According to an influential research program in cognitive science, philosophy, and linguistics, there is a deep, special connection between generics and pernicious aspects of social cognition such as stereotyping. Specifically, generics are thought to exacerbate our propensity to essentialize, lead us to overgeneralize based on scarce evidence, and lead to other epistemically dubious patterns of inference. Recently, however, several studies have put empirical and theoretical pressure on some of the main tenets of this research program. The goal of this paper is to bring these results together in a comprehensive narrative and systematically evaluate their impact on the hypothesis that generics have a uniquely problematic effect on our social and cognitive capacities.

1. Introduction

The last few decades have witnessed an explosion of research on the cognitive science of generic language and thought, resulting in an increasing consensus that generics (e.g., sentences like "Brits have bad teeth" or "sharks are dangerous") are in many ways *special*. For example, generics are acquired earlier in linguistic development than overt quantifiers such as "some" and "all", even though their truth-conditional profile is notoriously hard to capture (Gelman et al., 2015; Gelman & Raman, 2003; Graham et al., 2016; Hollander et al., 2002), and the generic operator *gen* is unpronounced in most, if not all, natural languages (Dahl, 1985; Krifka et al., 1995; Leslie, 2008).[3]

There's another way in which generics are thought to be special: the way in which they're connected to some of our most problematic forms of social reasoning. Many theories in linguistics, psychology, and philosophy posit a deep and unique connection between generics, on the one hand, and stereotypes and prejudice, on the other (see, e.g., Anderson et al., 2012; Antony, 2022; Appiah, 2018; Berio & Musholt, 2023; Bosse, 2022; Cappelen & Dever, 2019;

[1] This paper is forthcoming in *WIREs Cognitive Science*. Please cite the published version when available.

[2] For helpful feedback and discussion, we are grateful to Laurenz Casser, Joshua Knobe, Kevin Lande, Junhyo Lee, and two anonymous reviewers for *WIREs Cognitive Science*.

[3] See Kirkpatrick (2024) for critical discussion of acquisition data for generics; he argues that a trajectory similar to the one for generics arises for many other operators, including adverbial quantifiers. See Liebesman (2011) and Sterken (2015b) for discussion of whether the acquisition data and the fact that *gen* is unpronounced warrant the claim that generics are special.

Cimpian et al., 2010; Hammond & Cimpian, 2017; Haslanger, 2011; Hesni, 2024; Khemlani et al., 2012; 2011; Lerner & Leslie, 2016; Leslie, 2007, 2017; Rhodes et al., 2012; Wodak et al., 2015). As a result, many theorists have even proposed that avoiding using generics could be a successful stereotype mitigation technique (Leslie, 2017; Wodak & Leslie, 2017; Wodak, Leslie, & Rhodes, 2015).

The view that generics have a unique connection to various types of stereotypes has been particularly championed within the *Generics-as-Default* research program. Roughly, the Generics-as-Default hypothesis (henceforth: GaD) states that "generic sentences articulate cognitively fundamental, default generalizations, and that quantified statements, in contrast, articulate cognitively more sophisticated ones" (Leslie & Lerner, 2016; see also Gelman et al., 2015, 2019; Gelman, 2021; Johnston & Leslie, 2012; Khemlani, Leslie, & Glucksberg, 2009, 2011; Lerner & Leslie, 2016; Leslie, 2008, 2012, 2017; Prasada et al., 2013). Since stereotypes constitute some of the mind's default way of generalizing, and generics, according to GaD, give voice to and transmit the mind's defaults, it makes sense that there should be a deep connection between both. Generics—as Leslie puts it in her 2017 paper—contribute to the "Original Sin of Cognition" (p. 393).

An important project of the GaD program is to uncover the specific *properties* of this default system for generalization, which, by extension, will give us important insights into cognitive and linguistic forms of stereotyping (see, e.g., Lerner & Leslie, 2016, p. 407). As part of this effort, proponents of this program have proposed and investigated at least four ways in which generic generalizations *uniquely* (i.e., in contrast to sentences involving explicit quantification) contribute to stereotyping and prejudice:

> **Essentialism.** Generics convey that there is a principled and law-like—i.e., not merely accidental—relation between their subject and predicated property. Furthermore, this principled connection is due to the inherent or essential dispositions of the subject. In contrast, most overtly quantified sentences do not reliably convey inherent connections, even when they express a very high or reliable statistical relation. Hence the use of (even innocuous) generic language about a kind can lead interlocutors to adopt the (tacit) belief that the kind has an essence.

Strikingness. A generic 'Ks are F' is more likely to be accepted, especially at low prevalence levels, when the property F is striking or dangerous compared to when it is neutral or non-striking. Combined with the other hypotheses, this suggests that generics involving striking or dangerous properties are more likely to nudge us into asymmetrical inferences, overgeneralizing, and essentializing.

Overgeneralization. When presented with false universal statements of the form 'All Ks are F', adults and children often judge these statements to be true when the corresponding generic, 'Ks are F', is (judged to be) true. This suggests that our cognitive system interprets generalizations as generics by default. Since the processing of explicitly-quantified sentences has to override the default mode of generalizing, they are cognitively more costly, and as a result, people often default to a generic interpretation despite the presence of explicit quantification.

Asymmetry. Generics display a unique divergence between the conditions that are sufficient to accept their truth and the implications that are typically drawn from them. Suppose that, on average, most of us hold that at least *n*% of Ks have to be F to accept the generic 'Ks are F'. Yet when we accept a generic (say, on a speaker's authority) and predict what follows from that concerning the percentage of Ks that are F, most of us will tend to choose, on average, a percentage much higher than *n*. This divergence between acceptance and implications is thought to be unique to generics, compared to analogous statements with overt quantifiers and adverbs.

While each of these four properties has been put forward by proponents of the GaD program as a unique or characteristic property of generics (in contrast to explicitly quantified sentences), the view that generics are special in their pernicious effect on social cognition, in at least the four ways listed above, has been influential well beyond the GaD research program (see, e.g., Allaway et al., 2022; Appiah, 2018; Anderson, Haslanger, & Langton, 2012; Antony, 2022; Begby 2021; Berio & Musholt, 2023; Cappelen & Dever, 2019; Haslanger 2011; Rooij & Schulz, 2020). Given the vast influence of the view, and the enthusiastic and often uncritical acceptance of the main findings in its support, it is important to take a step back and think carefully about the landscape of the evidence. In fact, recently there has been a steady increase in empirical work that suggests that the relationship between generics and stereotyping, including Essentialism, Strikingness, Overgeneralization, and Asymmetry, may be more complex than it has been made out to be. Yet

this critical trend hasn't had much influence beyond a small circle of experts. The purpose of this essay is to highlight these findings, and incorporate them into a comprehensive narrative, focusing on how they affect our understanding of the four basic ways in which generics are thought to contribute to stereotyping and prejudice. For research on generics and social stereotypes to refrain from "spinning its own wheels" and avoid stagnation, these critical perspectives require attention.[4] We hope this overview will thereby help psychologists, cognitive scientists, and philosophers assess the overall state of that research program and open up new perspectives and questions.

2. Essentialism

2.1 The "Generics Essentialize" Hypothesis

Psychological essentialism is the hypothesis that there is "some hidden, nonobvious, and persistent property or underlying nature shared by members of that kind that causally grounds their common properties and dispositions" (Leslie, 2017, p. 405) and that makes members of the kind "the thing that they are" (Medin, 1989, p. 1476; see also Neufeld, 2022; Gelman, 2004). Belief in essence is often accompanied by a variety of further commitments, inferences, or expectations about the kind and its members. For example, essences are often represented as immutable, innately-determined, and internal rather than external. In addition, essentialized kinds are often taken to have sharp category boundaries, and to be discovered in nature, rather than invented (see Neufeld, 2022 for an overview). Thus, to essentialize a category or its members is often concomitant with treating the category in accordance with these further commitments, inferences, and expectations.

Haslanger (2011) has argued that generics about social kinds often cause us to incorrectly essentialize and form such concomitant beliefs: generics such as 'women are submissive' communicate that women's submissiveness is due to their nature—not to external or accidental social circumstances (see also Cimpian, Gelman, et al., 2010; Leslie, 2017; Prasada & Dillingham, 2006). If generics do lead us to essentialize, it is reasonable to expect that using generics would uniquely increase stereotyping and prejudice: generics would convey that there is a principled and law-like, hence inductively and modally powerful, connection between the subject and attributes.

[4] Park et al. (2023) observe that science has become less innovative over time, and that this stifles major advances or revisions.

Indeed, several studies suggest that the use of generic language increases essentializing for novel animals kinds (Cimpian et al., 2010; Gelman et al., 2010), familiar social kinds when generics match their prior knowledge (Cimpian & Erickson, 2012; Cimpian & Markman, 2011), and novel social kinds (Rhodes et al., 2018a; Rhodes et al., 2018b; Rhodes et al., 2012). For example, in Rhodes et al. (2012), adults and children (3 and 4 years) read a picture book about a novel social category of people called 'Zarpies.' The book described Zarpies using either generic language (e.g., 'Zarpies sleep in tall trees'), specifics with the category label (e.g., 'This Zarpie sleeps in tall trees'), or specifics without the category label (e.g., 'This one sleeps in tall trees'). After reading the book, participants were asked questions to measure the degree to which they essentialized Zarpies—e.g., whether an adopted Zarpie has the same characteristics as their biological Zarpie mother (essentializing) or their adoptive non-Zarpie mother (non-essentializing). The studies found that children and adults essentialized to a greater extent in the generics condition as compared to the specifics conditions. In an additional study of the same series, Rhodes et al. discovered that when parents were influenced to adopt essentialist views about Zarpies, the frequency of their generic statements about Zarpies to their children increased more than twofold. Based on these and other results, Lerner and Leslie (2016), along with several other theorists, conclude that "generics likely play a key role in the transmission of essentialist beliefs from parents to children" (Lerner & Leslie 2016, p. 11). Several theorists have since proposed to limit the use of generics about social groups, especially in conversations with children (Leslie, 2017; Rhodes et al., 2012; Wodak, Leslie, & Rhodes, 2015).

Recent work, however, suggests that the connection between generics and essentializing is less immediate and distinctive. To be sure, the issue is not whether generics *can* lead to essentializing, but rather whether it is generics *per se* that are the culprit in causing essentialist beliefs, and whether some of the original results may be driven by weaker or alternative commitments on the part of subjects than belief in an intrinsic essence. The rest of this section surveys work that suggests that the connection between generics and essentializing is more complicated than what's suggested by the 'Generics Essentialize' hypothesis. We first present studies that find a weaker or less distinctive link between generics and essentialism, and then consider studies which show that generics can systematically communicate extrinsic/external or structural information under certain conditions.

2.2 Empirical Evidence against the (Strong) Essentialism Hypothesis

According to the Generics Essentialize hypothesis, generics are special and distinctive in their capacity to cause essentialist thinking. However, Saul (2017) argues that it's not generics as such that are responsible for essentialization. Consider the following non-generic constructions:

(1) Many Asian men are abusive to women.
(2) My friend has a friend who's Asian and he's abusive to his wife.

Saul points out that these statements can lead to essentialization due to pragmatic considerations. In both cases, if the speaker didn't assume that there was some relevant *connection* between being Asian and abusiveness, they would hardly have mentioned the ethnicity. In addition, Saul argues that many studies that find a link between generics and essentialism have a flaw: they contrast generic constructions such as 'Ks are F' with 'This K is F' sentences, which leaves open the possibility that other constructions that signal high prevalence, such as 'Many/Usually Ks are F', are equally conducive to essentializing.

Inspired by these theoretical observations, Hoicka et al. (2021) investigated whether generics essentialize more than other quantified-sentences that signal high prevalence (e.g., 'many' or 'most'). They found that although generics and other high-prevalence quantified sentences led participants to *generalize*, surprisingly, neither generics, nor the other quantified sentences, led to an increase in *essentialism*. The results held true for both children and adults. That is, hearing statements about novel social kinds, such as 'Zarpies sleep in tall trees' and 'most Zarpies sleep in tall trees', led participants to believe that a new Zarpie also sleeps in tall trees. However, hearing either of these statements did not lead participants to judge that, for example, a Zarpie child will be more likely to sleep in tall trees when raised in a different environment. This suggests that generics lead to essentialization only under certain conditions, and that to the extent they lead to generalization of information, they do so in a way that doesn't substantially differ from that of other quantifiers signalling high proportions.

Similarly, Noyes and Keil (2019) present evidence that generics primarily communicate that their subject target category is a *kind* (as opposed to a mere *collection* of individuals or objects), but not that it is *essentialized*. What is the difference between being a kind and being essentialized? Kind representations encode a "reliable causal structure and high inductive potential" (p. 20354), whereas representing this kind as essentialized additionally requires that it is represented as

"naturally determined by internal causal properties" (p. 20354). Across multiple studies with an adult population, they found that it is only when the generic ascribes a biological property to a group (e.g., 'men grow beards') that participants infer an essentialist structure, but that neutral, cultural, or social generalizations (e.g., 'Jews believe eating pork is wrong') don't reliably communicate an essentialist category structure.

For example, in the first experiment, they presented participants either with generics ('Vawns are..') or specific sentences ('This Vawn is…'), and measured participants' responses to the kind vs. the essence dimensions. Their kind-measure included the statements:

1) Underneath superficial similarities and differences, all Vawns are basically the same.
2) Individual Vawns have very little in common. (Reverse-coded)
3) If someone tells you a fact about an individual Vawn, that fact is very likely true of other Vawns as well.
4) For some properties that Vawns have, it makes sense to say: "This person has that property because it is a Vawn."

Their essence-measure included the statements:

1) Vawns have internal or microscopic properties that cause their characteristic appearance and behavior.
2) The category Vawn was invented by people. (Reverse-coded)
3) The boundary between the category Vawn and non-Vawn is something decided by people. (Reverse-coded)
4) A Vawn can never change into a non-Vawn.

They found that generics had a large effect on the kind-measure, but a small one on the essence-measure. In their other experiments, they found that generics induce essentialism when they ascribe biological properties to a group, but not when they attribute social or neutral properties, and that describing a fake category as a biological vs. neutral/social kind leads to a comparable increase in likelihood for producing generics. In a different but related study, Bailey and Knobe (2023) found that reading biological-essentialist descriptions about fake categories does *not* increase the likelihood that participants endorse generic statements about that category.

An important lesson from the work surveyed in this sub-section is that the results of these types of studies hinge on our specific operationalizations of key notions such as 'essence', 'kind', or 'generalize'. For example, some theorists might take part of Noyes and Keil's 'kind' measures to correspond more closely to how essentialism and 'essence' is often characterized in the literature; their 'essence' measure, in turn, might be argued to correspond more closely to what others have used as an operationalization of 'natural kind'.[5] In addition, the research on structural representation and generics that we'll review in 2.2 suggests that generics do often trigger internalist explanations. Still, these studies suggest that the claim that generics lead to essentialism may have to be significantly refined, and that other grammatical constructions can lead to similar effects. We need more work on the specific conditions under which different kinds of generalizations do lead to essentialism, incl., work that is transparent and well-motivated in its choices of how to measure the generics-essentialism relationship.[6]

A further important lesson, also supported by the results in the next subsection (2.2), is that statements about the connection between generics and essentialism should not be understood as claims about the generic form *simpliciter*. The specific content of the generic matters crucially: e.g., whether the predicated property in 'Vawns are F' is biological as opposed to neutral or social, and plausibly also the strength of the generic generalisation itself is (e.g., 'squirrels are rodents' vs. 'squirrels like to eat street meat in Washington Square park'). Moreover, if the relevant contents can be communicated by other means (e.g., quantified sentences, implicatures, the given background knowledge), these may likewise result in essentialism. Future work should be careful to delineate factors such as form, content, and context in making claims about the relationship between generics and essentialism.[7]

[5] See Gelman, 2003, 2004; Haslam et al., 2000, 2004, 2014; Haslam & Levy, 2006; Mandalaywala, 2020; Neufeld, 2022; see Hoicka et al. pp. 4-5 for considerations distinguishing measures of essentializing vs. generalizing. See also pertinent literature in linguistics on well-established kinds (Carlson 1980; Chierchia 1998; Dayal 2004; Krifka et al. 1995; Mari et al. 2013; Liebesman & Sterken 2021) and ad hoc categories (Barsalou, 1983; Carston 2002; Mauri 2017; Mendia 2020; Wilson and Carston 2007), and their impact on the acceptability of generics.

[6] A broader lesson is, of course, a need for more meta-psychological discussions of operationalizations and measurements of psychological essentialism. For discussions of operationalization and measurement questions regarding psychological essentialism, see, e.g., Mandalaywala (2020), Rhodes and Mandalaywala (2017), Morgenroth (2021), Neufeld (2022), Haslam et al. (2004), Demoulin et al. (2006), Hoicka et al. (2021).

[7] Shahbazi et al. (2024) makes the important point that future work should also address different languages and cultures, which may be further sources of variability.

2.2 Generics and structural thinking

One way to explain an association or generalization like 'girls play with dolls' appeals to essential or inherent properties of girls—like an internal disposition to be caring and social. However, the studies we reviewed in the last subsection suggest that generics don't necessitate essentialist-internalist thinking. And in fact, there are other ways to explain generic associations or generalizations, including *structurally*. For example, a structural explanation of why girls play with dolls might go as follows: girls play with dolls because of aspects of the social context they're a part of—in particular, socio-historical factors that are 'external' to the girls in question, like that toy stores overwhelmingly stock dolls in the 'girls' section, can explain the generic. Structural explanations locate the source of the association between a kind and a property in the wider external structure in which the relevant kind or kind-members occupy a position. These structures are conceptualized as consisting of (social) positions or roles, and their relationships to one another. Such positions and relations are important as they indicate functions, properties, constraints, and enablements that kinds or kind-members possess, which serve as a rich basis for supposition, categorization, prediction, and inference (see Haslanger 2000, 2015; amongst others). Both internalist and structural representations can support stable and modally-robust, rather than merely accidental, associations or generalizations between kinds and properties. Yet they differ in whether these are supported by the internal properties of members of the kind (e.g., essences, dispositions or characteristic properties) or by external factors related to cultural, societal, or institutional structures.

Due to an apparently strong default tendency for essentialist and other forms of internalist thinking (Gelman 2003, 2004),[8] structural thinking has been thought of as challenging, unnatural or only available under limited circumstances (e.g., when an internalist explanation proves inadequate or a subject overrides their tendency for essentialist thinking). Recent work by Vasilyeva and colleagues, however, finds that it isn't too hard for subjects to interpret generics in

[8] Cimpian and Salomon (2014), Salomon and Cimpian (2014), and Cimpian (2015) propose that there is a tendency, termed the 'inherence heuristic', in which people explain patterns in association between a kind and a property in terms of inherent properties of members of the kind. Prasada and Dillingham (2006, 2009) argue that some kind-property associations communicated through generics are understood as inherent aspects of that kind. Gelman (2003) and Gelman et al. (2010) propose that receiving information expressed through generics leads to essentializing as described in the previous section. Leslie (2007, 2008) argues that some generics are made true just in case the properties ascribed are characteristic of members of the kind where what's characteristic is spelled out along roughly internalist lines. She also argues that generics are by default interpreted as expressing "generalizations that hold because of common, inherent features of the members of the kind" (2014, p. 217), drawing on Gelman. See Nadya Vasilyeva, Alison Gopnik, and Tania Lombrozo (2018) for an informative overview of 'internalist' views of generics and generic representations.

a structural-externalist, rather than an internalist, way. In their (2018) study, Vasilyeva, Gopnik, and Lombrozo tested whether children and adults were able to represent information conveyed via generics in structural terms. In the study, a novel category property association (children playing a fictional ball game) was introduced in different framing conditions designed to induce either internalist or structuralist construals. Participants were then given different tasks in which they were asked to explain the associations (where their responses were coded as *internalist*, *structural*, or *other*) or judge explanations that were offered to them. The study showed that, given appropriate cues, generics can robustly and easily trigger structural explanations in 5-6 year olds and adults. The study did also show that, in the absence of structuralist cues, most participants produced internalist responses. However, this tendency decreased with age. In sum, while structural representations of associations conveyed via generics occurred in both adults and children, the frequency and ease with which they were generated increased with age.

In a further study, Vasilyeva and Lombrozo (2020) reproduced these results in adults using sample sentences that concerned a fictional immigrant group and more realistic properties. The 2020 study replicated findings from the 2018 study, namely that both internalist and structural representations of associations conveyed via generics could be generated in adults. It also found that in the absence of structural cues, internalist representations were more frequent. This suggests that "while structural interpretations can be induced with minimal cues, they are not the default or habitual interpretation of generic claims attributing properties to social groups, at least under the conditions assessed by our task" (pp. 7–8).

Building on the insight that generics can convey structural explanations, recent work rethinks previous proposals about how to mitigate negative forms of essentializing. For example, Saul (2017) argues that because generics are well-equipped to express social critique, a general prohibition on using them is misguided. Similarly, Vasilyeva and Ayala-López (2019) argue that insights from the structural approach may be used to elicit, from particular uses of generics, structural explanations, i.e., as a form of intervention against internalist thinking and thereby counter the associated discriminatory behavior (see also Hesni, 2024; McKeever & Sterken, 2021; Ritchie, 2019 for discussion). In support of this, Vasilyeva, Gopnik, and Lombrozo (2020) conducted a study with 3-8 year old children in which they tested whether interventions can break the link between generics and essentializing by providing context that elicits structural representations. They found that with appropriate cues, children were able to represent relevant

associations in structural terms and that this significantly decreased the extent to which they formed essentialist expectations.

In sum, theoretical and empirical work on generics and structural explanations suggests that generics are not completely tied to internalist explanations. To be sure, we have seen that, in the absence of specific cues to elicit structural representations, internalist representations seem to prevail, especially in younger children. Still, the result that structural representations can be easily and systematically elicited shows that the link between generics and internalist representations and essentializing can be reliably broken. Of course, more work needs to be done to discover what other cues and prior information increase the likelihood that generics are interpreted as expressing or eliciting structural explanations. At this point, it is wise to wait for such results before the practical significance of any proposed interventions on the use of generics can be seriously assessed.

3. Striking Property Generics

According to many defenders of the GaD hypothesis, generics will lead to problematic stereotypes in a particularly ampliative way when the attributed properties are "striking, horrific, or appalling" (Leslie, 2017, p. 410). Leslie (2017) explains this in the following way: if generics voice generalizations that are cognitively fundamental, and we can't "wait around to see what percentage of tigers eat people before drawing a general conclusion" (p. 396), it follows that we should be particularly prone to accept a generic as true when it attributes a dangerous property, and act accordingly:

> "It is not hard to imagine the evolutionary benefits of such a disposition, since the costs of undergeneralizing such information are potentially huge. Our ancestors were far better off jumping to conclusions, as it were, than taking the time to judiciously determine the precise likelihood of their being eaten" (Leslie, 2017, p. 396).[9]

According to this view, striking property generics are extremely tolerant to exceptions and are accepted even if only a small proportion or a few exemplars of the category possess the target property. Applied to stereotyping and prejudice, this means we will be more likely accept a

[9] A further illustrative quote: "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, 2017, p. 410).

generic as true of a social group when the predicated property is dangerous or otherwise negative, or accept a generic as true of the whole kind when only very few members of the category possess the relevant striking property. Empirical evidence for this view is presented by Cimpian et al. (2010), who found that adults accept generics at a lower prevalence when the attributed property is dangerous or distinctive. In addition, Prasada et al. (2013) found that most subjects (between 70-77%) accept striking property generics such as 'Sharks attack bathers' as true (yet note that they still had lower acceptance rates than other types of generics they tested, e.g., minority and majority characteristic generics).

Against the Striking-Property hypothesis, recent empirical evidence found that this effect doesn't neatly generalize across various major kinds of generics. Tasimi et al. (2017) found that although using striking property predicates leads to an increased inclination to accept generics for artifacts and non-human animal categories, striking predicates do *not* increase subjects' willingness to accept generics for human categories. These initial results suggest that Leslie (2008; 2017) and Cimpian et al.'s (2010) original claims about striking property generics should be substantially qualified—at least, it is still an open question whether the basic effect generalizes uniformly across different categories and type of generics of interest, including human social categories.

Lazaridou-Chatzigoga et al. (2019b) conducted experiments which also put pressure on a central socially relevant aspect of the Striking-Property Hypothesis, namely, that although we easily accept generics about striking *and uncommon* features, we then go on to attribute the property very broadly to the whole category and its instances. In Experiment 1, Lazaridou-Chatzigoga et al. tested whether the strikingness of a property used in a generic to describe a category makes a difference to how willing children and adults were to apply that same property to novel members of that same category (they call this a 'property extension' task). When testing children, they showed them puppets of a novel kind (e.g., 'Glippets') and described them with a neutral or striking generic (e.g., 'These are glippets. Glippets love to play with toys / fire'). These descriptions were chosen to feature 'child-friendly' items with high arousal and average valence as a proxy for strikingness based on affective ratings from Warriner et al. (2013). Next, they showed children a novel instance of the same kind, and asked whether that instance will have the property attributed in the generic, too (e.g., 'Does this glippet love to play with toys / fire?').[10] As expected, the researchers found that both children's and adults' projection of properties to novel instances was sensitive to the type of property used in the generic. But contrary to the

[10] The adult version was analogous, but used pictures and text instead of puppets.

Striking-Property hypothesis, both children and adults were *less* willing to extend a property to a new instance when the kind was described with a striking property, rather than a neutral, generic. The findings suggest that "the fact that the property is striking and less common leads participants to be conservative and extend striking properties less than neutral properties" (Lazaridou-Chatzigoga et al., 2019b, p. 5).

In their second study, Lazaridou-Chatzigoga et al. also tested whether property extension tasks based on generics are less sensitive to exceptions for striking property generics compared to neutral generics. As in Experiment 1, two instances of a novel creature were introduced and a property was attributed using a neutral or striking property generic. Next, they presented participants with either one or three *exceptions* to the target generalization (e.g., "But not this one. This glippet doesn't love to play with toys/fire"). Then, a fourth/sixth instance of the kind was presented along with a property extension question ("Does this glippet love to play with toys/fire?"). Following Leslie (2008) and Cimpian et al. (2010), it was expected that property extensions based on striking property generics would be less sensitive to exceptions than extensions based on neutral generics. But again, this was *not* what they found. While subjects did extend properties in the face of exceptions in both the neutral and striking property conditions, they found that subjects were substantially *less* tolerant of exceptions for striking property generics than for neutral generics.

What is the upshot of these results? We are not denying that there is some truth to Leslie's general observation that, from an evolutionary perspective, we are and should be especially sensitive to striking or dangerous properties. Yet these new studies challenge her specific proposal that the basic way in which our cognition manifests that sensitivity is by systematically (and irrationally) exaggerating judgments about the *likelihood* that members of target groups have the relevant striking properties. The acceptance of the relevant striking property generics (to the extent that it is observed), and corresponding behavioral effects (e.g., cautious behaviour), can instead result from differences in *negative expected utility* estimates.[11] Those estimates can result from a combination of roughly *accurate* judgments about the distribution of dangerous properties with the very high negative valence assigned to those properties. For example, you might want to avoid swimming with sharks, even if you know that the probability of being attacked by one is very low, because the cost of being attacked is so high. Acting cautiously in the face of dangerous properties, then, need not be based on a tendency to over*generalize*.

[11] Generally, expected utility is the weighted average of all possible outcomes of a decision, with each outcome weighted by its probability and the utility (or value) it provides.

Indeed, this simple expected-utility model helps us make sense of the basic results of Lazaridou-Chatzigoga et al., according to which participants use the kind of property attributed in a generic to guide their likelihood estimates. By drawing on their background information as speakers, subjects can infer that, in general, a lower prevalence of the predicated property is needed for people to accept a striking compared to a neutral property generic. So as recipients of (novel) striking property generics, subjects can then rationally infer that the striking property in question is likely not as widely-distributed compared to a similar situation involving a neutral property. Finally, this model opens an intriguing way of trying to understand why the acceptance of striking property generics is dependent on the *type* of category the generic is about, as shown in Tasimi et al.'s study. By refusing to engage with certain artifacts and non-human animals, we often incur a cost. However, our cost-estimates may be systematically higher when it comes to refusing to engage with various human or social groups. This suggests that we should take seriously the possibility that, contrary to what is sometimes suggested by advocates of the GaD program, reasoning with generics (boundedly) rationality, and in particular doesn't lead us to systematically over-generalize properties to categories, even when they're striking.

4. The Overgeneralization Effect

One of the major claims of the GaD program is that all explicitly-quantified sentences, including universally-quantified ones, are represented as generics *by default*, especially in tasks with high cognitive load. This is because more sophisticated modes of generalization have to 'override' our default mode, which is the one invoked by generics, according to GaD. This explains how and why generics are connected to stereotyping: If we have a tendency to understand and recall quantified sentences as generics *by default*, then information of the form 'some/at least one/25%/65%/most/many/etc. of Ks are F' is often processed and remembered as generics. In many cases, this would have the effect of both inflating the perceived prevalence of Ks that are F and (via essentializing) supporting the belief that an inherent relation sustains that association.

Many studies have been presented as supporting this hypothesis (Brandone et al., 2015; Gelman et al., 2016, 2019; Hollander et al., 2002; Leslie et al., 2011; Leslie & Gelman, 2012; Mannheim et al., 2010; Meyer et al., 2011; Ralethe & Buys, 2022; Tardif et al., 2012). Most of these studies focus on the quantifier "all". For example, when children and adults are presented with strictly false universal sentences such as '*all* ducks lay eggs', they judge them to be true when the corresponding generic sentence is also judged as true. This suggests that subjects tacitly interpret

the universal sentences as generics. This effect has been coined as the *Generic Overgeneralization Effect* (henceforth, *GOG*).[12]

In an influential study in support of the GOG effect, Leslie et al. (2011) found that even adult participants accept false universal generalizations as true when the generic counterpart would be true.[13] In their first experiment, participants were asked to judge the truth of generic and universal statements. Most participants judged the 'all'-statements to be true, although they were false, and only their corresponding generic version was true. The numbers are impressive: 78% of so-called "majority characteristic" statements (e.g., 'all tigers have stripes') and 51% of so-called "minority characteristic" statements (e.g., 'all ducks lay eggs') were judged to be true. In contrast, when the property in question was accidental and the corresponding generic false, as in 'all books are paperbacks', only 13% of all-statements were judged as true. This suggests that both children and adults overgeneralize "from the truth of a generic to the truth of the corresponding universal statement" (Leslie et al., 2011, p. 17).

To rule out that the GOG effect is due to ignorance of exceptions—e.g., in the case of misjudging 'all ducks lay eggs' as true, due to ignorance of the fact that not all ducks lay eggs—Leslie et al. conducted a test to determine if participants knew that, say, 'male ducks lay eggs' is false. Although most participants knew about those kinds of non-trivial exceptions, the GOG effect didn't disappear. Yet when the knowledge test was presented before evaluating the 'all'-statement, the percentage of incorrectly evaluated 'all' minority statements dropped from 51% to 19%. To rule out that the GOG effect was due to a subkind interpretation (e.g., interpreting the 'all'-statement as 'all duck *breeds* lay eggs'), participants were asked to provide paraphrases of the statements that they were supposed to judge. The experimenters found that only 1% of paraphrases appealed to subtypes. Finally, to rule out that the effects were due to quantifier domain restriction (e.g., interpreting the 'all'-statement as 'all *female* ducks lay eggs'), the participants were provided, before the targets, with background contexts which consisted of population estimates such as "Suppose the following is true: there are 431 million ducks in the world. Do you agree with the following: all ducks lay eggs?" The idea is that the context nudges participants to set the domain of quantification to every duck in the world. Still, there was a GOG effect: participants still accepted 60% of majority universal statements as true (a drop of

[12] Curiously, and possibly in tension with GOG, many experiments also found that generics are interpreted as if universally-quantified (Reuter et al., 2023; Cimpian et al., 2010; Brandone et al., 2014).
[13] Hollander, Gelman and Star (2002) did find an Overgeneralization Effect for children, but not for adults.

18%), and accepted 30% of minority statements as true (a drop of 21%). In each case, then, introducing the control reduced the GOG effect, but didn't make it disappear.

One concern, at this point, is that even if we take the GOG effect at face value, the experiments just focus on whether generalizations with the explicit quantifier "all" seem to behave by default as generics. Accordingly, we would still need evidence for the claim that we interpret various other explicitly quantified sentences by default as generics.[14]

But even if we set this issue aside, Lazaridou-Chatzigoga, Katsos, and Stockall argue in several papers that there are arguably better ways to explain the basic GOG effects (Lazaridou-Chatzigoga, 2019; Lazaridou-Chatzigoga et al., 2015; Lazaridou-Chatzigoga, Katsos, et al., 2019; Lazaridou-Chatzigoga, Stockall, et al., 2019; Lazaridou-Chatzigoga et al., 2023; Lazaridou-Chatzigoga & Stockall, 2013; see also Castroviejo et al., 2023). First, they point out that although Leslie et al. controlled for *each* of the possible alternative explanations individually, they didn't control for *all* of them—that is, they didn't rule out that all the explanations *taken together* fully account for the GOG effect. Secondly, they argue that Leslie et al. didn't control for the atypical behaviour of the quantifier 'all'. Work in linguistics suggests that 'all' is often used hyperbolically as meaning something akin to 'very many' (Claridge, 2010), and that 'all', in contrast to 'every' and 'each', is ambiguous between distributive and collective interpretations (Beghelli & Stowell, 1997). Accordingly, sentences such as 'all ducks lay eggs' could be interpreted as meaning that ducks *collectively* lay eggs, or that (very many) individual ducks lay eggs.[15]

Lazaridou-Chatzigoga et al.'s main criticism, however, is that Leslie and colleagues don't pay sufficient attention to work on the semantics and pragmatics of quantifier domain restriction,

[14] There are several other open theoretical questions that the claim that we interpret all explicitly quantified sentences by default as generics gives rise to: if this is the case, it is not clear why generic *sentences* would be special for the transmission of stereotypes, essentialism, and so on, as is repeatedly emphasized by proponents of the GaD program (cf. Wodak et al., 2015; Leslie, 2014). Similarly, if we, by default, represent all explicitly-quantified information as generic, it is not clear why generics are immune to counter-examples (expressed via sentences with overt quantifiers). As Peters, Kraus, and Braganza (2022) point out, if GaD was correct, we should represent counterexamples (e.g., "Some women are good at math") in terms of generics, too: "If people had an automatic tendency to generalize that is context-independent, then they should also readily generalize from some counter-stereotypical individuals to the corresponding kind of them. Since this does not commonly happen, the impact of generalization bias is likely influenced by psychological, contextual, and social factors" (p. 7).

[15] Note that collective readings are harder to get for accidental majority generics and their corresponding 'all'-statements, e.g., artifact kind 'all'-statements like 'all cars have radios' or 'all books are paperback'.

and their dismissal of that alternative explanation of the GOG effect isn't convincing. To illustrate the basic phenomenon, consider:

(3) There was rhubarb in the pie. Everyone developed a rash. (von Fintel, 1994)

When interpreting the second sentence of (3), it is clear that the domain of 'everyone' is restricted to the members of the specific pie-eating situation. Since these kinds of domain restrictions are very common in natural language, it is important to seriously attempt to appeal to them when trying to explain the GOG effect.

Lazaridou-Chatzigoga et al. propose that one way to test the quantifier domain restriction hypothesis is to introduce contexts that make the relevant domains for the restriction salient. The idea is that by using (a) a neutral context that doesn't interact with the content of the critical statement, (b) a contradictory context that makes salient an exception and rules out an unrestricted universal quantification, and (c) a supportive context that makes the generality of the property in question salient, and various types of universal quantifiers that show different sensitivity to quantifier domain restriction, the GOG and Quantifier Domain Restriction (QDR) hypotheses can be pitted against each other.

To illustrate, consider the following examples of context and target sentences (cf. Lazaridou-Chatzigoga et al., 2017):

> *Neutral Context:* Linton Zoo is home to three tigers, Tibor, Baginda and Kaytlin, whose playful games visitors love to watch and photograph.
>
> *Contradictory Context:* Linton Zoo is home to three tigers, Tibor, Baginda and Kaytlin, whose fur is all white due to a recessive gene that controls coat color.
>
> *Supportive Context:* Linton Zoo is home to three tigers, Tibor, Baginda and Kaytlin, whose black and orange coats visitors love to photograph.
>
> *Target Sentences:* All tigers have stripes/All the tigers have stripes/Each tiger has stripes.

The universal quantifier 'all' allows, but doesn't require, quantifier domain restriction to be tied to a specific or salient situation, while the quantifiers 'all the' and 'each' require discourse-linked domain restriction in virtue of their semantics (Partee, 1995; Pesetsky, 1987). Crucially, the QDR and GOG hypotheses make different predictions about people's response patterns. GOG predicts that, for each of the three conditions, the three quantifiers should lead to roughly the same level of error judgments. QDR predicts that people shouldn't be biased by the corresponding generic interpretation, but instead base their judgements on the specific quantifier and context. For example, contradictory contexts should lead to very low acceptance rates for 'all the' and 'each', but higher acceptance rates for 'all', since the latter doesn't require domain restriction to that specific situation and tends to be more open to hyperbolic or kind interpretations. Finally, supportive contexts should lead to high acceptance rates for the sensitive quantifiers, and neutral contexts establish a baseline to observe difference in effects for different quantifier types.

Lazaridou-Chatzigoga, Stockall, et al. (2019) tested these two hypotheses, implementing their proposed strategy in multiple languages, and found that their results supported the predictions made by the QDR hypothesis. That is, contrary to the predictions made by the GOG hypothesis, the error rates between 'all', 'each', and 'all the' differ as expected given the sensitivity of each quantifier to domain-restriction (see also Lazaridou-Chatzigoga, Katsos, et al., 2019a). This criticism of the GOG hypothesis suggests a general lesson: quantificational language is a well-studied phenomenon in theoretical and experimental linguistics, and it is crucial for proponents of the GOG hypothesis to seriously engage with this body of work when selecting controls for their studies and making substantial background assumptions about how various quantifiers work.

5. The asymmetry of generics

According to the Asymmetry hypothesis, generics display a distinctive divergence between their acceptance conditions and their implications: while relatively low prevalence values can be sufficient to accept a generic as true, we tend to then infer a substantially higher prevalence level when taking a generic as a premise: e.g., the proportion of mosquitos that needs to carry the West Nile virus for us to accept as true the generic 'Mosquitos carry the West Nile virus' tends to be substantially lower than the proportion that we would tend to infer from when, say, a perceived expert communicates that generics to us. As before, the surprising claim is not just that generics display this asymmetry between acceptance and implication conditions, but that they

display it *uniquely*. Other constructions, such as overtly quantified sentences with "most" or "many", allegedly do not lead to this inference pattern.[16]

In arguably the most influential paper supporting Asymmetry, Cimpian, Brandone, and Gelman (2010) present a series of studies which they interpret as supporting the view that "the prevalence level that is sufficient to judge a generic sentence as true is indeed significantly lower than the prevalence level implied by that very same sentence and [...] that this asymmetry is peculiar to generic statements and does not extend to sentences with quantified noun phrases as subjects" (p. 1453). In the Truth Condition (TC) task, they presented participants with information about prevalence of the properties true of novel categories, and measured the average level that leads to acceptance of the associated generic sentence. For example, depending on assignment, participants were told that 10%, 30%, 50%, 70%, or 90% of "lorches" have purple feathers, and were then asked whether the generic "Lorches have purple feathers" is true. In the Implied Prevalence (IP) task, Cimpian et al. assessed the average prevalence level implied by a generic. Participants were presented with a generic such as "Lorches have purple feathers", and were then asked what percentage of lorches they estimate to have purple feathers. To test whether the hypothesized asymmetry holds uniquely for generics, they presented a different group of participants with analogous tasks but using 'most' sentences. In support of Asymmetry, Cimpian et al. (2010) found that the average prevalence that led participants to accept generics was significantly lower than the average prevalence implied by them (~69 vs. ~96%). No such difference was found for the quantifier 'most', where the average was almost identical in both conditions (~77% vs. 78%). Using a very similar paradigm, later studies tested and confirmed Asymmetry for children (Brandone et al., 2015) and generics about the social domain (Cella et al., 2022).

Asymmetry has been widely accepted in philosophy, psychology, and linguistics, and invoked as a significant contributor to the emergence of social and political biases (Bian & Cimpian, 2017; Cappelen & Dever, 2019; Gelman & Roberts, 2017; Haslanger, 2011; Haward et al., 2018; Leslie,

---

[16] According to Leslie and Lerner, Asymmetry is directly predicted by GaD:

> "If the generics-as-defaults hypothesis is correct, then the acceptance of a generic "Xs are Y" should be closely tied to a person's default inferences: a person's tendency to infer that an arbitrary member of kind X has property Y. Empirical work by Cimpian, Brandone, and Gelman (2010) and Khemlani, Leslie, and Glucksberg (2012) confirms this prediction. Cimpian et al. found that people regularly inferred from statements of the form "Xs are Y" that almost every X is Y, despite believing that "Xs are Y" could be true even if only a minority of Xs were Y" (Lerner & Leslie, 2016, p. 10).

While it is rarely spelled out in which ways GaD really *predicts* Asymmetry, the thought seems to be, we take it, that the tendency to infer beyond one's evidential basis is in some sense irrational and unreflective, and must therefore reflect some sort of default-system (see also Peters et al., 2022).

2017; McKeever & Sterken, 2021; Peters, Krauss, & Braganza, 2022; Rooij & Schulz, 2020; Rosola & Cella, 2020; Sterken, 2015a). It is easy to understand the enthusiasm, for Asymmetry would, if true, likely have a substantial impact on the dynamics of information exchange in social networks. To illustrate why, note that "experts" should be willing to accept a generic like "sharks are dangerous" when, according to their observations and beliefs, the proportion of sharks that are dangerous meets the minimum threshold for its truth. Now suppose that experts then communicate that information to (non-expert) "recipients" using precisely that generic sentence. When the recipients accept it, they will, due to Asymmetry, be disposed to infer that the proportion of sharks which are dangerous is significantly greater than the proportion which supported the original introduction of the generic. Over time and iterations of this asymmetry effect, we are likely to observe a substantial gap between the actual and perceived prevalence levels for the properties of many kinds, leading to inaccurate stereotypes and biases.

However, in a recent series of studies, Reuter, Neufeld, and Del Pinal (2023, 2024) presented evidence against Asymmetry. Those studies maintain various aspects of Cimpian's original study, but modify what they argue are important methodological shortcomings. First, Reuter et al. compared generics with a wider range of quantifiers and adverbs (e.g., "usually", "typically"). The latter are particularly important since, according to many linguists, they are arguably closest in meaning to the unpronounced generic quantifier *Gen*. At any rate, a diverse and inclusive comparison class of generalizing sentences is required to seriously test whether Asymmetry is really unique to generics. Second, some design choices made by Cimpian et al (2010), and repeated in subsequent studies, might have artificially erased a potential asymmetry effect for 'most'. In the TC conditions (for 'most' or generic sentences), participants were presented with a prevalence level at one of five predetermined interval values (10%, 30%, 50%, 70%, 90%), and then asked whether the 'most' or generic sentence is true at that level. To determine the prevalence level required for the truth of a "most" or generic sentence, Cimpian et al. then averaged over all prevalence values at which participants accept as true the target generalizing sentence. However, this method doesn't capture the prevalence level that is sufficient for the truth of a generalizing sentence. Consider: if there are 100 rhinos, 'at least one rhino has four legs' is true at all prevalence interval values between 1% and 100%. If, as we would expect, participants answered 'true' for each value, and we averaged over all of those conditions, we would get a result of roughly 50%. But this clearly doesn't capture the truth conditions of the "at least one" statement, which is closer to the notion of the minimal prevalence level sufficient for truth. To address this shortcoming, Reuter and colleagues provided participants with a

continuous percentage scale which they used to set both the minimum prevalence for acceptance of a generalizing sentence, and the implied prevalence level which they would infer given the truth of a corresponding generalizing sentence.

Using this minimally modified experimental design, Reuter et al. found that all tested (non-exact) generalizing sentences displayed a significant asymmetry effect. Specifically, all other tested quantifiers and adverbs—i.e. "most", "some", "typically", "usually"—also exhibited an asymmetry effect. This seems to align with straightforward linguistic intuitions: if a speaker asks someone to bring them "some lemons", and she agrees, the speaker would normally expect that she would bring the speaker a bunch, and not exactly one or two—even if the speaker was prepared to agree that bringing, say, exactly one is strictly speaking compatible with the instructions.

It might be tempting to think that, ultimately, this result extends a core proposal of the GaD program, namely, that generalizing language systematically exacerbates our social biases and stereotypes. This conclusion would be premature. The TCs tasks in all these studies ask participants at what prevalence level they would accept that some generalizing sentence is true. They do not ask participants whether, given a range of options, they would also choose to use that target sentence to communicate some specific prevalence information. Yet generalized asymmetry effects would only exacerbate biases if, from a range of options, speakers would often choose to communicate information about a given prevalence level *n* using a generalizing sentence that has an acceptance level of roughly around *n* but an implied prevalence level that is significantly higher. It is easy to illustrate why that might well be a rare occurrence, at least in cooperative situations. Suppose that 52% of lorches have blue feathers: although that prevalence level might be strictly sufficient for you to accept as true "most lorches have blue feathers", but it doesn't commit you to use precisely that sentence to communicate that information to a recipient (e.g., a sentence using "roughly half" could be much better choice).

Reuter, Neufeld, and Del Pinal (2024) tested whether communicators behave in this way---i.e., when choosing which expression to actually use are systematically sensitive not just to their acceptance conditions but also to the implications likely drawn by their typical recipients. They found that, in cooperative contexts, speakers tend to choose, from a range of options, generalizing sentences whose implied prevalences best match the target prevalence values which they wish to communicate. This aligns with other results from the Rational Speech Act program,

which suggest that speakers tend to choose which sentence to assert based on models that are sensitive to the implications that pragmatic listeners are likely to draw (Degen, 2023; Goodman and Frank, 2016). Indeed, Reuter et al. also found that, in strategic contexts where their goals and those of their interlocutors fail to align, speakers do systematically choose sentences which increase the difference between actual and likely implied prevalence values, when that discrepancy is likely to serve their goals. This suggests that the asymmetry effects associated with generalizing sentences can, given certain conditions and motivations, be exploited to introduce biases into a discourse (Banaji et al., 2021; Kunda, 1990)—but generics are not unique in lending themselves to these kinds of strategic, exploitative uses.

Conclusion

The last few decades have seen an explosion of research on the connection between generics and pernicious forms of social cognition. This project, inspired by and still deeply connected with the GaD-program, has led to important insights and advanced our understanding of generic language and thought. Yet the enthusiasm surrounding the program and its intriguing and often surprising findings has led many of us to overlook nuances, conflicting empirical and theoretical perspectives, and important insights from adjacent disciplines. Our goal in this critical review is to present, in a more systematic way than has been done so far, the main empirical and theoretical challenges to the view that generics have a uniquely negative effect on our social judgments and related cognitive capacities.

Zooming out, our review points to a substantial contrast between the GaD program and related work on generics and social cognition, and the various critical perspectives and findings that we have presented. The GaD approach treats generic cognition and communication as a species of passive cognition, in the sense that it characterizes cognitive processes involving generics as part of a default/automatic, irrational, unconscious, unreflective cognitive system.[17] Generic language and thought is then expected to inherit many problematic features of this system, including the propensity to exacerbate problematic social biases and stereotypes. In contrast, several of the critical perspectives and results we reviewed seem to indicate that generic representations do not neatly fit this framework. Instead, generic cognition and communication may involve more active, reflective, and rational processes. From this perspective, generic *and* quantificational thought and communication in general involve various choice points: in particular, the need to

[17] Correspondingly, it has sometimes aligned itself with the dual-process theory of cognition, according to which there is a fast, unreflective, and automatic system ("system 1") that houses generics, and a slow, deliberate, rational system ("system 2") that houses explicitly-quantified sentences (Almotahari, 2024; Leslie, 2007; see also Kirkpatrick, 2024; Lazaridou et al., 2023; vaan Rooij & Schulz, 2020; Sorensen, 2013).

estimate prevalences, propensities and how those relate to truth-conditions which need not be directly about those values, the need to engage in conscious reflection and deliberation on how counterexamples and exceptions affect generalizations and interact with loose speech, the need to fix domain of quantification and resolve for ambiguity, semantic interactions and underdetermination, as well as to make strategic language choices given motivational, practical and aspects of the context including interlocutors likely pragmatic uptakes (cf. Bosse, 2021, 2022; Kirkpatrick 2024; Lazaridou-Chatzigoga et al. 2013, 2015, 2019a, 2019b, 2019c, 2023; Lee & Nguyen, 2022; McKeever & Sterken, 2021; Nguyen, 2019; Nickel, 2016; Plunkett et al., 2023; Saul, 2017; Sterken, 2015b; Tessler & Goodman, 2019). The GaD theory and its extensions to the interface between generics and social cognition, at least in its unqualified version, doesn't seem well equipped to accommodate these active elements.

Still, our point here isn't that the GaD program and its offshoots should be abandoned. Proponents of the GaD may be able to explain away on independent grounds or refine their theories in response to most of the challenges we have presented (see, e.g., Almotahari, 2024). Nor is our point that generics never contribute to problematic forms of stereotyping. The point is rather that their role is arguably less unique and direct---and much more complex and dependent on factors shared by other kinds of non-exact generalizations---than has thus far been acknowledged. We hope our review can help guide and inspire new empirical work that continues to explore and refine the various fascinating hypotheses which have come out of that research program.

References

Allaway, E., Taneja, N., Leslie, S.-J., Sap, M. (2022) Towards countering essentialism through social bias reasoning. NLP4PosImp Workshop, EMNLP. https://doi.org/10.48550/arXiv.2303.16173

Almotahari, M. (2024). Generic cognition: A neglected source of context sensitivity. *Mind & Language*. https://doi.org/10.1111/mila.12491

Anderson, L., Haslanger, S., & Langton, R. (2012). Language and race. *Routledge Companion to the Philosophy of Language, London, Routledge*, 753-767.

Antony, L. (2022). The Importance of Being Partial: The Constructive Role of Bias in Human Life. *The Amherst Lecture in Philosophy*, *15*, 1–19.

Appiah, K. A. (2018). *The Lies that Bind: Rethinking Identity*. Profile Books.

Bailey, A. H., & Knobe, J. (2023). Biological essentialism correlates with (but doesn't cause?)

intergroup bias. *Personality and Social Psychology Bulletin*, 01461672231158095.

Banaji, M. R., Fiske, S. T., & Massey, D. S. (2021). Systemic racism: Individuals and interactions, institutions and society. *Cognitive Research: Principles and Implications*, *6*, 1–21.

Barsalou, L. W. (1983). Ad hoc categories. *Memory & cognition*, *11*, 211-227.

Begby, E. (2021). *Prejudice: A Study in Non-Ideal Epistemology*. Oxford: Oxford University Press.

Beghelli, F., & Stowell, T. (1997). Distributivity and negation: The syntax of each and every. *Ways of Scope Taking*, 71–107.

Berio, L., & Musholt, K. (2023). How language shapes our minds: On the relationship between generics, stereotypes and social norms. *Mind & Language*, *38*(4), 944–961. https://doi.org/10.1111/mila.12449

Bian, L., & Cimpian, A. (2017). Are stereotypes accurate? A perspective from the cognitive science of concepts. *Behavioral and Brain Sciences*, *40*, e3. https://doi.org/10.1017/S0140525X15002307

Bosse, A. (2021). Generics: Some (Non) Specifics. *Synthese*, *5–6*, 14383–14401. https://doi.org/10.1007/s11229-021-03426-8

Bosse, A. (2022). Generics and Stereotyping. *Inquiry,* 1-17, https://doi.org/10.1080/0020174X.2022.2074879

Brandone, A. C., Gelman, S. A., & Hedglen, J. (2015). Children's developing intuitions about the truth conditions and implications of novel generics versus quantified statements. *Cognitive Science*, *39*(4), 711–738. https://doi.org/10.1111/cogs.12176

Cappelen, H., & Dever, J. (2019). *Bad Language*. Oxford University Press.

Carlson, G., 1980, Reference to Kinds in English. Garland, New York & London.

Carston, R. (2002). Linguistic meaning, communicated meaning and cognitive pragmatics. *Mind & Language*, *17*(1-2), 127-148.

Castroviejo, E., Hernández-Conde, J. V., Lazaridou-Chatzigoga, D., Ponciano, M., & Vicente, A. (2023). Are generics defaults? A study on the interpretation of generics and universals in 3 age-groups of Spanish-speaking individuals. *Language Learning and Development*, *19*(3), 275-302. https://doi.org/10.1080/15475441.2022.2071715

Cella, F., Marchak, K. A., Bianchi, C., & Gelman, S. A. (2022). Generic Language for Social and Animal Kinds: An Examination of the Asymmetry Between Acceptance and Inferences. *Cognitive Science*, *46*(12), e13209. https://doi.org/10.1111/cogs.13209

Chierchia, G., 1998, Reference to kinds across languages. Natural Language Semantics 6(4): 339-405.

Cimpian, A. (2015). The inherence heuristic: Generating everyday explanations. *Emerging Trends in*

the Social and Behavioral Sciences: An Interdisciplinary, Searchable, and Linkable Resource, 1–15.

Cimpian, A., Brandone, A. C., & Gelman, S. A. (2010). Generic statements require little evidence for acceptance but have powerful implications. *Cognitive Science*, *34*(8), 1452–1482. https://doi.org/10.1111/j.1551-6709.2010.01126.x

Cimpian, A., & Erickson, L. C. (2012). The effect of generic statements on children's causal attributions: Questions of mechanism. *Developmental Psychology*, *48*(1), 159.

Cimpian, A., Gelman, S. A., & Brandone, A. C. (2010). Theory-based considerations influence the interpretation of generic sentences. *Language and Cognitive Processes*, *25*(2), 261–276. https://doi.org/10.1080/01690960903025227

Cimpian, A., & Markman, E. M. (2011). The generic/nongeneric distinction influences how children interpret new information about social others. *Child development*, *82*(2), 471-492.

Cimpian, A., & Salomon, E. (2014). The inherence heuristic: An intuitive means of making sense of the world, and a potential precursor to psychological essentialism. *Behavioral and Brain Sciences*, *37*(5), 461–527.

Claridge, C. (2010). *Hyperbole in English: A corpus-based study of exaggeration*. Cambridge University Press.

Dahl, Ö. (1985). Tense and aspect systems. Oxford: Blackwell.

Dayal, V., 2004, Number Marking and (In)definiteness in Kind Terms. Linguistics and Philosophy 27.4, 393-450.

Degen, J. (2023). The rational speech act framework. *Annual Review of Linguistics*, 9, 519-540.

Demoulin, S., Leyens, J.-P., & Yzerbyt, V. (2006). Lay Theories of Essentialism. *Group Processes & Intergroup Relations*, *9*(1), 25–42. https://doi.org/10.1177/1368430206059856

Gelman, S. A. (2003). *The Essential Child: Origins of Essentialism in Everyday Thought*. Oxford University Press.

Gelman, S. A. (2004). Psychological essentialism in children. *Trends in Cognitive Sciences*, *8*(9), 404–409. https://doi.org/10.1016/j.tics.2004.07.001

Gelman, S. A. (2021). Generics in society. *Language in Society*, *50*(4), 517-532.

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. https://doi.org/10.1080/23273798.2014.931591

Gelman, S. A., Leslie, S.-J., Gelman, R., & Leslie, A. (2019). Do Children Recall Numbers as Generic? A Strong Test of the Generics-As-Default Hypothesis. *Language Learning and Development*, *15*(3), 217–231. https://doi.org/10.1080/15475441.2019.1571418

Gelman, S. A., & Raman, L. (2003). Preschool Children Use Linguistic Form Class and

Pragmatic Cues to Interpret Generics. *Child Development*, *74*(1), 308–325. https://doi.org/10.1111/1467-8624.00537

Gelman, S. A., & Roberts, S. O. (2017). How language shapes the cultural inheritance of categories. *Proceedings of the National Academy of Sciences*, *114*(30), 7900–7907. https://doi.org/10.1073/pnas.1621073114

Gelman, S. A., Sánchez Tapia, I., & Leslie, S.-J. (2016). Memory for generic and quantified sentences in Spanish-speaking children and adults. *Journal of Child Language*, *43*(6), 1231–1244. https://doi.org/10.1017/S0305000915000483

Gelman, S. A., Ware, E. A., & Kleinberg, F. (2010). Effects of generic language on category content and structure. *Cognitive psychology*, *61*(3), 273-301.

Goldfarb, D., Lagattuta, K. H., Kramer, H. J., Kennedy, K., & Tashjian, S. M. (2017). When your kind cannot live here: How generic language and criminal sanctions shape social categorization. *Psychological Science, 28*(11), 1597–1609. https://doi.org/10.1177/0956797617714827

Goodman, N. D. & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Science*, 20(11), 818-829.

Graham, S. A., Gelman, S. A., & Clarke, J. (2016). Generics license 30-month-olds' inferences about the atypical properties of novel kinds. *Developmental Psychology*, *52*(9), 1353–1362. https://doi.org/10.1037/dev0000183

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), 607–614. https://doi.org/10.1037/xge0000297

Haslam, N., Holland, E., & Karasawa, M. (2014). Essentialism and Entitativity Across Cultures. In *Culture and Group Processes* (pp. 17–37). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199985463.003.0002

Haslam, N., & Levy, S. R. (2006). Essentialist Beliefs About Homosexuality: Structure and Implications for Prejudice. *Personality and Social Psychology Bulletin*, *32*(4), 471–485. https://doi.org/10.1177/0146167205276516

Haslam, N., Rothschild, L., & Ernst, D. (2000). Essentialist beliefs about social categories. *The British Journal of Social Psychology*, *39 (Pt 1)*, 113–127. https://doi.org/10.1348/014466600164363

Haslam, N., Rothschild, L., Ernst, D., Yzerbyt, V., Judd, C. M., & Corneille, O. (2004). Essentialism and Entitativity: Structures of Beliefs about the Ontology of Social

Categories. *The Psychology of Group Perception: Perceived Variability, Entitativity, and Essentialism.*, 61–78. https://doi.org/10.4324/9780203644973

Haslanger, S. (2000). Gender and race: (What) are they? (What) do we want them to be? *Noûs*, *34*(1), 31–55. https://doi.org/10.1111/0029-4624.00201

Haslanger, S. (2011). Ideology, generics, and common ground. In *Feminist Metaphysics: Explorations in the Ontology of Sex, Gender and the Self* (pp. 179–207). Springer Netherlands. http://www.springerlink.com/index/10.1007/978-90-481-3783-1_11

Haward, P., Wagner, L., Carey, S., & Prasada, S. (2018). The development of principled connections and kind representations. *Cognition*, *176*, 255-268. https://doi.org/10.1016/j.cognition.2018.02.001

Hesni, S. (2024). Generics and social justice. *Philosophical Studies*. https://doi.org/10.1007/s11098-023-02064-9

Hoicka, E., Saul, J., Prouten, E., Whitehead, L., & Sterken, R. (2021). Language Signaling High Proportions and Generics Lead to Generalizing, but Not Essentializing, for Novel Social Kinds. *Cognitive Science*, *45*(11), e13051. https://doi.org/10.1111/cogs.13051

Hollander, M. A., Gelman, S. A., & Star, J. (2002). Children's interpretation of generic noun phrases. *Developmental Psychology*, *38*(6), 883–894. https://doi.org/10.1037//0012-1649.38.6.883

Johnston, M., & Leslie, S.-J. (2012). Concepts, Analysis, Generics and the Canberra Plan. *Philosophical Perspectives*, *26*(1), 113–171. https://doi.org/10.1111/phpe.12015

Khemlani, S., Leslie, S. J., & Glucksberg, S. (2009). Generics, prevalence, and default inferences. *Proceedings of the 31st annual cognitive science society*, 443-448.

Khemlani, S., Leslie, S.-J., & Glucksberg, S. (2012). Inferences about members of kinds: The generics hypothesis. *Language and Cognitive Processes*, *27*(6), 887–900. https://doi.org/10.1080/01690965.2011.601900

Kirkpatrick, J. R. (2024). The acquisition of generics. *Mind & Language*, 1–26. https://doi.org/10.1111/mila.12499

Krifka, M., Pelletier, F. J., Carlson, G., Ter Meulen, A., Chierchia, G., & Link, G. (1995). Genericity: An Introduction. In Greg N. Carlson & Francis Jeffry Pelletier (eds.), *The Generic Book*. University of Chicago Press. pp. 1-124.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*(3), 480–498. https://doi.org/10.1037/0033-2909.108.3.480

Lazaridou-Chatzigoga, D. (2019). Genericity. In C. Cummins & N. Katsos (Eds.), *The Oxford Handbook of Experimental Semantics and Pragmatics* (p. 0). Oxford University Press.

https://doi.org/10.1093/oxfordhb/9780198791768.013.12

Lazaridou-Chatzigoga, D., Katsos, N., & Stockall, L. (2015). Genericity is easy? Formal and experimental perspectives. *Ratio*, *28*(4), 470–494.

Lazaridou-Chatzigoga, D., Katsos, N., & Stockall, L. (2019a). *Experimental evidence on genericity and universal quantification in Greek and English*. 171–182.

Lazaridou-Chatzigoga, D., Katsos, N., & Stockall, L. (2019b). Generalizing About Striking Properties: Do Glippets Love to Play With Fire? Frontiers in Psychology, 10. https://doi.org/10.3389/fpsyg.2019.01971

Lazaridou-Chatzigoga, D., & Stockall, L. (2013). *Genericity, exceptions and domain restriction: Experimental evidence from comparison with universals*. *17*, 325–343.

Lazaridou-Chatzigoga, D., Stockall, L., & Katsos, N. (2017). A new look at the 'Generic Overgeneralisation' effect. *Inquiry*, *66*(9), 1655–1681. https://doi.org/10.1080/0020174X.2017.1285993

Lazaridou-Chatzigoga, D., Stockall, L., & Katsos, N. (2019c). Contextualising generic and universal generalisations: Quantifier domain restriction and the generic overgeneralisation effect. *Journal of Semantics*, *36*(4), 617–664.

Lazaridou-Chatzigoga, D., Stockall, L., & Katsos, N. (2023). A new look at the 'Generic Overgeneralisation' effect. *Inquiry*, *66*(9), 1655–1681. https://doi.org/10.1080/0020174X.2017.1285993

Lee, J., & Nguyen, A. (2022). What's positive and negative about generics: A constrained indexical approach. *Philosophical Studies*, *179*(5). https://doi.org/10.1007/s11098-021-01727-9

Lerner, A., & Leslie, S.-J. (2016). Generics and Experimental Philosophy. In *A Companion to Experimental Philosophy* (pp. 404–416). John Wiley & Sons, Ltd. https://doi.org/10.1002/9781118661666.ch28

Leslie, S.-J. (2007). Generics and the Structure of the Mind. *Philosophical Perspectives*, *21*(1), 375–403. https://doi.org/10.1111/j.1520-8583.2007.00138.x

Leslie, S.-J. (2008). Generics: Cognition and Acquisition. *Philosophical Review*, *117*(1), 1–47. https://doi.org/10.1215/00318108-2007-023

Leslie, S.-J. (2012). Generics Articulate Default Generalizations. *Recherches Linguistiques de Vincennes*, *41*, Article 41. https://doi.org/10.4000/rlv.2048

Leslie, S. J. (2012b). Generics. In G. Russell & D. G. Fara (Eds.), *Routledge companion to philosophy of language* (pp. 355–367). Routledge.

Leslie, S.-J. (2014). "Carving Up the Social World with Generics." In *Oxford Studies in Experimental*

Philosophy, edited by Joshua Knobe, Tania Lombrozo, and Shaun Nichols, 208–231. Oxford: Oxford University Press.

Leslie, S.-J. (2017). The Original Sin of Cognition. *Journal of Philosophy*, *114*(8), 393–421. https://doi.org/10.5840/jphil2017114828

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. https://doi.org/10.1016/j.cogpsych.2011.12.001

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.

Leslie, S.-J., & Lerner, A. (2016). Generic Generalizations. *The Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/entries/generics/

Liebesman, D., & Sterken, R. K. (2021). Generics and the metaphysics of kinds. *Philosophy Compass*, *16*(7), e12754.

Mandalaywala, T. (2020). Does essentialism lead to racial prejudice? It is not so Black and White. In *Adv Child Dev Behav.* (pp. 195–245). https://doi.org/10.1016/bs.acdb.2020.05.007

Mannheim, B., Gelman, S. A., Escalante, C., Huayhua, M., & Puma, R. (2010). A developmental analysis of generic nouns in Southern Peruvian Quechua. *Language Learning and Development*, *7*(1), 1–23.

Mari, A., Beyssade, Del Prete, F. (ed.) 2013, Genericity. OUP.

Mauri, C., 2017, Building and lnterpreting Ad Hoc Categories: A Linguistic Analysis. In J. Blochowiak, C.Grisot, St. Durrleman, Ch. Laenzlinger (eds) Formal Models in the Study of Language. Springer, 297-326.

McKeever, M., & Sterken, R. (2021). Social and Political Aspects of Generic Language and Speech. In *The Routledge Handbook of Social and Political Philosophy of Language*. Routledge.

Mendia, J.A. (2020). Reference to ad hoc kinds. *Linguistics and Philosophy 43*, 589–631.

Medin, D. L. (1989). Concepts and Conceptual Structure. *American Psychologist*, *44*(12), 1469–1481. https://doi.org/10.1037/0003-066X.44.12.1469

Meyer, M., Gelman, S. A., & Stilwell, S. M. (2011). *Generics are a cognitive default: Evidence from sentence processing. 33*(33).

Neufeld, E. (2022). Psychological essentialism and the structure of concepts. *Philosophy Compass*, *17*(5), e12823. https://doi.org/10.1111/phc3.12823

Nguyen, A. (2019). The radical account of bare plural generics. *Philosophical Studies*, 1–29. https://doi.org/10.1007/s11098-019-01254-8

Nickel, B. (2016). *Between logic and the world: An integrated theory of generics*. Oxford University Press.

Nickel, B. (2017). Generics. In *A Companion to the Philosophy of Language* (pp. 437–462). John Wiley & Sons, Ltd. https://doi.org/10.1002/9781118972090.ch18

Noyes, A., & Keil, F. C. (2019). Generics designate kinds but not always essences. *Proceedings of the National Academy of Sciences*, *116*(41), 20354–20359. https://doi.org/10.1073/PNAS.1900105116

Park, M., Leahey, E., & Funk, R. J. (2023). Papers and patents are becoming less disruptive over time. *Nature*, *613*(7942), 138–144.

Partee, B. H. (1995). Quantificational structures and compositionality. In *Quantification in natural languages* (pp. 541–601). Springer.

Pesetsky, D. (1987). Wh-in-situ: Movement and unselective binding. *The Representation of (in) Definiteness*, *98*, 98–129.

Peters, U., Krauss, A., & Braganza, O. (2022). Generalization bias in science. *Cognitive science*, *46*(9), e13188.

Plunkett, D., Sterken, R. K., & Sundell, T. (2023). Generics and Metalinguistic Negotiation. *Synthese*, *201*(50), 1–46. https://doi.org/10.1007/s11229-022-03862-0

Prasada, S., & Dillingham, E. M. (2006). Principled and statistical connections in common sense conception. *Cognition*, *99*(1), 73–112. https://doi.org/10.1016/J.COGNITION.2005.01.003

Prasada, S., & Dillingham, E. M. (2009). Representation of Principled Connections: A Window Onto the Formal Aspect of Common Sense Conception. *Cognitive Science*, *33*(3), 401–448. https://doi.org/10.1111/j.1551-6709.2009.01018.x

Prasada, S., Khemlani, S., Leslie, S.-J., & Glucksberg, S. (2013). Conceptual distinctions amongst generics. *Cognition*, *126*(3), 405–422.

Ralethe, S., & Buys, J. (2022). Generic Overgeneralization in Pre-trained Language Models. In N. Calzolari, C.-R. Huang, H. Kim, J. Pustejovsky, L. Wanner, K.-S. Choi, P.-M. Ryu, H.-H. Chen, L. Donatelli, H. Ji, S. Kurohashi, P. Paggio, N. Xue, S. Kim, Y. Hahm, Z. He, T. K. Lee, E. Santus, F. Bond, & S.-H. Na (Eds.), *Proceedings of the 29th International Conference on Computational Linguistics* (pp. 3187–3196). International Committee on Computational Linguistics. https://aclanthology.org/2022.coling-1.282

Reuter, K., Neufeld, E., & Del Pinal, G. (forthcoming). Generic and Quantified Generalizations: Asymmetry Effects and Strategic Communication. *Cognition.*

Reuter, K., Neufeld, E., & Del Pinal, G. (2023). Asymmetry Effects in Generic and Quantified Generalizations. *Proceedings of the 45th Meeting of the Cognitive Science Society 45*, 1-6. https://escholarship.org/uc/item/4zk277m7

Rhodes, M., Leslie, S. J., Bianchi, L., & Chalik, L. (2018a). The role of generic language in the early development of social categorization. *Child Development*, *89*(1), 148-155.

Rhodes, M., Leslie, S. J., Saunders, K., Dunham, Y., & Cimpian, A. (2018b). How does social essentialism affect the development of inter-group relations?. *Developmental science*, *21*(1), e12509.a

Rhodes, M., Leslie, S. J., & Tworek, C. M. (2012). Cultural transmission of social essentialism. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(34), 13526–13531. https://doi.org/10.1073/pnas.1208951109

Rhodes, M., & Mandalaywala, T. (2017). The development and developmental consequences of social essentialism. *Wiley Interdisciplinary Reviews: Cognitive Science*, *8*(4), e1437. https://doi.org/10.1002/wcs.1437

Ritchie, K. (2019). Should We Use Racial and Gender Generics? *Thought: A Journal of Philosophy*, *8*(1), 33–41. https://doi.org/10.1002/tht3.402

Rooij, R. van, & Schulz, K. (2020). Generics and typicality: A bounded rationality approach. *Linguistics and Philosophy*, *43*(1), 83–117. https://doi.org/10.1007/s10988-019-09265-8

Rosola, M., & Cella, F. (2020). Generics and Epistemic Injustice. *Ethical Theory and Moral Practice*, *23*(5), 739–754. https://doi.org/10.1007/s10677-020-10095-y

Salomon, E., & Cimpian, A. (2014). The Inherence Heuristic as a Source of Essentialist Thought. *Personality and Social Psychology Bulletin*, *40*(10), 1297–1315. https://doi.org/10.1177/0146167214541659

Saul, J. (2017). Are generics especially pernicious? *Inquiry*, *66*(9), 1682–1699. https://doi.org/10.1080/0020174X.2017.1285995

Shahbazi, G., Hossein, S., Mandalaywala, T.M., Borhani, K., Davoodi, T. (2024). When do generics lead to essentialism: Developmental evidence from Iran. *Infant and Child Development*. https://doi.org/10.1002/icd.2538 Sorensen, R. (2012). The sorites and the generic overgeneralization effect. *Analysis*, *72*(3), 444-449.

Sterken, R. K. (2015a). Generics, Content and Cognitive Bias. *Analytic Philosophy*, *56*(1), 75–93. https://doi.org/10.1111/phib.12056

Sterken, R. K. (2015b). Generics in Context. *Philosophers' Imprint*, *15*(i), 1–30.

Tardif, T., Gelman, S. A., Fu, X., & Zhu, L. (2012). Acquisition of generic noun phrases in Chinese: Learning about lions without an '-s.' *Journal of Child Language*, *39*(1), 130–161.

Tasimi, A., Gelman, S. A., Cimpian, A., & Knobe, J. (2017). Differences in the Evaluation of Generic Statements About Human and Non-Human Categories. *Cognitive Science*, *41*(7), 1934–1957. https://doi.org/10.1111/cogs.12440

Tessler, M. H., & Goodman, N. D. (2019). The language of generalization. *Psychological Review*, *126*(3), 395–436. https://doi.org/10.1037/rev0000142

Vasilyeva, N., & Ayala-López, S. (2019). Structural Thinking and Epistemic Injustice. In B. R. Sherman & S. Goguen (Eds.), *Overcoming Epistemic Injustice: Social and Psychological Perspectives*. Rowman & Littlefield International.

Vasilyeva, N., Gopnik, A., & Lombrozo, T. (2018). The development of structural thinking about social categories. *Developmental Psychology, 54*(9), 1735–1744. https://doi.org/10.1037/dev0000555

Vasilyeva, N., Gopnik, A., & Lombrozo, T. (2020). When generic language does not promote psychological essentialism. In *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*.

Vasilyeva, N., & Lombrozo, T. (2020). Structural thinking about social categories: Evidence from formal explanations, generics, and generalization. *Cognition*, *204*, 104383. https://doi.org/10.1016/j.cognition.2020.104383

Von Fintel, K.-U. (1994). Restrictions on quantifier domains. *Doctoral Dissertations Available from Proquest*, 1–295.

Warriner, A.B., Kuperman, V., & Brysbaert. M. (2013) Norms of valence, arousal, and dominance for 13,915 English lemmas. Behaviour research methods, 45, 1191-1207

Wilson, D., & Carston, R. (2007). *A unitary approach to lexical pragmatics: Relevance, inference and ad hoc concepts* (pp. 230-260).

Wodak, D., & Leslie, S. J. (2017). The mark of the plural: Generic generalizations and race. In *The Routledge companion to the philosophy of race* (pp. 277-289). Routledge.

Wodak, D., Leslie, S. J., & Rhodes, M. (2015). What a loaded generalization: Generics and social cognition. *Philosophy Compass*, *10*(9), 625-635. https://doi.org/10.1111/phc3.12250