



Generics: some (non) specifics

Anne Bosse¹ 

Received: 25 February 2021 / Accepted: 21 September 2021
© The Author(s) 2021

Abstract

This paper is about an underappreciated aspect of generics: their non-specificity. Many uses of generics, utterances like ‘Seagulls swoop down to steal food’, express non-specific generalisations which do not specify their quantificational force or flavour. I consider whether this non-specificity arises as a by-product of context-sensitivity or semantic incompleteness but argue instead that generics semantically express non-specific generalisations by default as a result of quantifying existentially over more specific ones.

Keywords Generics · Genericity · Non-specificity · Propositional pluralism · Context-sensitivity · Semantic incompleteness

1 Introduction

Finding and articulating the conditions under which generics, sentences like *Bus drivers are grumpy* or *Bats have good hearing*, are true or false has been the principal undertaking in the generics literature. This is because it’s important, but also because it’s difficult. It’s important because understanding the truth conditions of generics tells us, at least in part, what they mean. And we should want to know what they mean. An active research programme on generics has recently drawn attention to their cognitive, political, and moral significance.¹ However, determining the truth-conditions of generics has been difficult because the facts that make them true vary radically. Generics can be true in virtue of facts about what most members of a kind are like,

¹For instance, see Sarah Jane Leslie’s work on generics as articulating cognitively basic default generalisations (2007, 2008, 2012) and that they act as vehicles of stereotyping (2015, 2017). On the latter, also see Anderson et al (2012), McConnell-Ginet (2012), Haslanger (2011, 2014), Saul (2017), Ritchie (2019), and Lemeire (2020).

This article belongs to the topical collection "Indeterminacy and Underdetermination", edited by Mark Bowker and Maria Baghramian.

✉ Anne Bosse
annebosse@gmail.com

¹ University of Cambridge, Peterhouse, Trumpington Street, Cambridge CB2 1RD, UK

what they are disposed to be like, or even what they should be like, amongst others.² In this paper, I propose a simple explanation of this fact: generics can be made true by different kinds of facts because the generalisations they express are non-specific.

I start by articulating an assumption made implicitly in much of the literature on generics: that generics express generalisations that are specific with respect to quantificational force and flavour. I present a counterexample to this assumption: a type of utterance which does not express any one specific, non-generic generalisation.

1.1 Non-specificity

All generalisations are non-specific in the sense that they are not about particular members of a category. For example, *All dogs bark* or *Most dogs bark* are not about any particular dogs.³ However, in this paper, I will argue that generics are non-specific in a further sense: they do not specify the quantificational force or flavour of the link between members of a kind and the property in question. While *Most dogs bark* specifies a prevalence link between dogs and barking, the corresponding generic, *Dogs bark*, does not further specify the nature of the link it posits between dogs and barking.

I consider whether this non-specificity arises as a by-product of context-sensitivity or semantic incompleteness, but argue instead that the propositions generics like *Dogs bark* express are themselves weak generalisations that can be made true by facts of different quantificational strengths and flavours, like the fact that most dogs bark, but also other facts, such as that barking is characteristic of dogs, or that more dogs than cats bark.

2 The specificity assumption and a counterexample

This paper is about generics, sentences like (1)–(3):⁴

- (1) Rocking chairs are wooden
- (2) Elks have antlers
- (3) Brazilians are good football players

Although the literature on generics is young, there are plenty of theories about their truth conditions. According to one popular approach, generics are universal generalisations in disguise, quantifying over *normal* members of a kind.⁵ In contrast, Ariel Cohen (1999a, 1999b) argues that generics are about what members of a kind are

² For discussion of this feature, see Leslie (2008), Sterken (2015a), and Nguyen (2019).

³ Thanks to an anonymous reviewer for this journal for pointing out this further sense in which generics are non-specific.

⁴ I focus my attention on bare-plural generics in English for the purposes of this paper. See Greenberg (2004, 2007) for discussion of non-bare plural generics.

⁵ Theories of this type, albeit with various refinements, have been advocated by Dahl (1975), Delgrande (1987), Asher and Morreau (1995), Eckardt (2000), Greenberg (2004, 2007), Asher and Pelletier (2012), and Nickel (2016).

likely to be like.⁶ A different approach is taken by Sarah-Jane Leslie (2008, 2012), who argues that generics are made true by different kinds of facts, depending on the property they ascribe. Some generics are made true by facts about what's true of most members of a kind, others by facts about what's characteristic of the kind, and yet others by facts about what properties members of the kind are disposed to have.

One assumption these views generally make is that generics express fine-grained generalisations: generalisations about what's normal, what's likely to be the case, what's characteristic, or what members of a kind are disposed to be like.⁷ On this assumption, generics are not that different from other types of sentences that express generalisations, at least not with regards to their meaning. Even though generics may lack expressions like 'all' or 'usually' at the surface level, the generalisations they express are not different in kind from those expressed by overtly quantified generalisations: they are generalisations with a particular flavour and quantificational force. So, whatever generalisations generics actually express, those generalisations are *either* about what normal members of a kind are like, what members of a kind are likely to be like, and so on. It's this assumption I want to question by providing an example of generics use which is not specific in that manner.

2.1 Seagulls

My friend Aoife and I are discussing things she doesn't like about Aberdeen. So far, Aoife has talked about the greyness, lack of sunlight, and bad city planning. Then she says: 'And what I hate most is the seagulls because we can't eat anything outside. Seagulls swoop down to steal food.' The last part of Aoife's remark is a generic, but what type of generalisation does it express?

(4) Seagulls swoop down to steal food

- (a) Most seagulls swoop down to steal food
- (b) All normal seagulls swoop down to steal food
- (c) All seagulls are disposed to swoop down to steal food
- (d) More seagulls than other types of birds swoop down to steal food
- (e) Swooping down to steal food is characteristic of seagulls

I will call (4a)–(4e) 'candidate generalisations'. The assumption made in the views I have described so far is that Aoife's utterance expresses one candidate generalisation or some other parsing of comparable specificity. However, I don't think that it is at all obvious that this is the case and that only the facts that would make true *one* of (4a)–(4e) are what could make true Aoife's utterance of (4).

First, decide on whichever parsing of (4) strikes you as the most plausible. Now, imagine that that parsing was false, but that another of the candidate generalisations

⁶ According to Cohen, generics receive one of two readings. They can have absolute prevalence readings in which case they are true iff the likelihood of a given relevant member of the kind having the property in question is greater than 0.5. Alternatively, generics can receive a comparative reading, in which case they are true iff the probability of a given member of the kind having some probability is higher than for a member of a salient comparison kind.

⁷ The matter is a little bit more complicated for normalcy approaches in that they tend to include a contextual parameter in the truth-conditions of generics. I will discuss contextualist accounts in Sect. 4.1 of this paper.

was true. Would you want to say that Aoife's utterance was, strictly speaking, false? If not, then it seems like the assumption is unwarranted.

Second, sentences like (4) are often used to convey rough generalisations. Assuming that assertions are typically used to express beliefs, is it plausible to think that the belief Aoife expresses, and hence her communicative intention, must be as fine-grained as the belief she would have expressed had she uttered one of (4a)–(4e)? It doesn't seem so to me. We often use generics to articulate rough generalisations, in this case about seagulls and some behaviour that we think they *in some general sense* exhibit. Using generics in this way just doesn't require believing generalisations like (4a)–(4e).

This example is not an anomaly but an instance of a wider phenomenon. Much of our use of generics is not specifically about normal members of a kind, what we take to be the probability of a member of a kind having some property compared to some comparison class, and so on. Instead, we use generics to articulate generalisations that are not specific in this way. This paper is about understanding the non-specificity at hand.

One initial thought might be that the context in which Aoife utters (4) matters. Recently, accounts according to which sentences like (4) can convey *different* generalisations depending on the context in which they are uttered have gathered momentum. Given that those espousing views of this kind tend to place an emphasis on the speaker's intentions, they might also have something to say about contexts in which speakers appear to lack the fine-grained intentions that would be required to make one of (4a)–(4e) be what Aoife's utterance expressed.

And indeed, both Rachel Sterken (2015a) and Anthony Nguyen (2019), who have emphasised flexibility in the use of generics across contexts, have discussed what they call 'underspecified' or 'indeterminate' uses of generics. I will first outline their views and the resources that come with them, before considering how they can be brought to bear on non-specific uses of generics.

3 Token variability

Sterken and Nguyen distinguish between two different ways in which the generalisations conveyed by generic tokens appear to vary across contexts of utterance: in their quantificational force and flavour.⁸

3.1 Quantificational force variability

Nguyen uses the following example:⁹

(5) Lottery tickets are losers

John needs money and asks Sally whether she thinks he should buy lottery tickets. In response, Sally says: 'No, lottery tickets are losers'. Nguyen argues that in the

⁸ While Sterken thinks that the variability concerns the truth-conditions of the generic tokens, Nguyen thinks it occurs in the implicatures conveyed by them. I explain the difference in Sect. 4.

⁹ Example originally given in Sterken (2015a, p. 18) and developed by Nguyen (2019, p. 1305).

context of Sally's utterance, the generalisation the generic token conveys is true iff *most* lottery tickets are losers. Contrast this with a different token of the same generic type. Donald believes that the lottery is a big scam. He signs a contract, stating that he'll give any lottery winner a billion US dollars. His friend, who thinks this was a stupid idea, objects. Donald defends himself by saying 'Lottery tickets are losers'. According to Nguyen, the generalisation conveyed by this generic token is true iff all lottery tickets are losers and consequently, even for one and the same generic sentence type, generalisations conveyed by different tokens of it can have different truth-conditions.

This phenomenon is not specific to (5). Take the following example:

(6) Pizzas are in the freezer section

(6) can be used to make claims of varying quantificational force about pizzas and their location. Imagine two supermarkets, one in which all the pizzas are in the freezer section and another in which most are in the freezer section, but there are also some in the fridge section. A customer comes into the first supermarket and asks a shop assistant where to find pizzas. The shop assistant might utter (6) to make the claim that all the pizzas are in the freezer section. However, if the customer instead went into the second supermarket, a shop assistant there could use (6) to make a much weaker claim, such that *most* or even just *many* pizzas are in the freezer section.

3.2 Flavour variability

Sterken gives examples in which the truth-conditions of utterances of one and the same generic type appear to vary, not in their quantificational force, but the *type* of generalisation they make.¹⁰ I call this variability 'flavour variability'.

(7) Horses wear horseshoes

Sterken argues that (7)¹¹ is true in some contexts and false in others. In the wild, no horses wear horseshoes. They are not available and horses don't need them. However, it is plausible that, as a result of the size of the domesticated horse population, most horses wear horseshoes.

Tokens of (7) asserted in a discussion of modern horse care are true. However, tokens of (7) asserted contexts in which questions of evolutionary biology are salient, say in a TV programme about evolutionary adaptations, are, according to Sterken, false.¹² She argues that this is because the truth conditions of tokens of (7) vary across contexts. Tokens of (7) asserted in a modern horse care context, say in a conversation between horse breeders, are true iff most horses wear horseshoes. Tokens of (7) asserted in an evolutionary biology context are true iff horses naturally wear horseshoes. Since the

¹⁰ Sterken (2015a, pp. 6–7).

¹¹ Sterken actually uses the much more widely discussed example of *Dobermanns have pointy ears* introduced by Nickel (2008). It is possibly no longer the case that most dobermanns have their ears clipped. Because the example as presented relies on some prevalence claim being true, I use (7) instead. The considerations are analogous.

¹² Nguyen (2019, p. 1306) gives some further examples of this kind.

former is the case and the latter isn't, tokens of (7) are, Sterken argues, true in some contexts and false in others.¹³

I do think that the examples given by Sterken and Nguyen pick up on a variability in use, though whether that variability is truth-conditional or of a different kind is a question I want to postpone until Sects. 7 and 8. Even though Sterken and Nguyen seem to assume that many of the tokens they use in their examples are specific enough to trigger the changes in truth-value judgements in question, both also discuss non-specific uses of generics. As they take non-specificity to arise out of token variability, we need to first understand how they explain the latter to understand how they make sense of the former.

4 Explaining token variability

Sterken and Nguyen explain token variability by positing contextualism and semantic incompleteness respectively.

4.1 Contextualism

Sterken uses the observations about token variability to motivate an account, according to which generics are context-sensitive. She argues that an unpronounced quantifier generics are often thought to contain, Gen, is context-sensitive.¹⁴ Following Jeffrey King (2013), Sterken argues that Gen is a supplementive. Supplementives are context-sensitive expressions, which, in addition to their context-invariant meaning and the context of utterance, require that speakers intend to convey a certain content and that this intention is accessible to an attentive, reasonable hearer who knows the common ground of the conversation. For example, the semantic values of demonstratives like 'that' or pronouns like 'her' depend on the intentions of the speaker, in addition to their context-invariant meaning and other features of the context.

As the intentions of speakers vary across contexts of utterance, both the flavour and quantificational force of Gen do too. For example, speaker intentions, according to Sterken, can make it such that a token of (4), *Seagulls swoop down to steal food*, expresses the claim that most seagulls swoop down to steal food in one context, whereas another utterance of (4) expresses the claim that seagulls are more likely to swoop down to steal food than other salient kinds of bird. According to contextualism then, the truth conditions of generics vary across tokens, depending on certain features of the contexts of utterance, notably the speaker's intentions. An approach I consider next explains variability not in terms of the generics being context-sensitive, but in terms of them being semantically incomplete.

¹³ Another kind of notable type of flavour variability is discussed in Leslie (2015) who considers cases in which utterances of a generic can have a normative flavour in one context and a descriptive flavour in another.

¹⁴ Sterken's account is just one way of implementing contextualism about generics. For example, most of the normalcy accounts I mentioned in Sect. 2 include a contextual parameter.

4.2 Semantic incompleteness

Nguyen (2019) proposes that generics are semantically incomplete. He utilises Kent Bach's notion of propositional radicals to articulate his proposal. A propositional radical is a structured proposition missing at least one of its constituents. Consequently, the would-be proposition has a gap. By themselves, the sentences that express the propositional radicals lack truth-values. However, they can be completed by speaker intentions. For example, 'It is raining' doesn't, in itself, express a full proposition, but can be used by a speaker to make the claim that it is raining in São Paulo, as long as the speaker has the relevant communicative intention.¹⁵ According to Bach, the completed content is not the semantic value of the utterance but is communicated pragmatically as an implicature.¹⁶

In contrast to contextualism, completed tokens do not have full propositions as their semantic value, but express them pragmatically. Whereas for contextualists, the semantic value of tokens varies across contexts of utterances, for Nguyen, the semantic content stays the same. Instead, it is pragmatically communicated content that varies depending on the speaker's intentions that complete propositional radicals in one way or another.

Applied to generics, the idea is that the generic types themselves do not have truth-values as they, in isolation, express propositional radicals which are not truth-apt. Nevertheless, generic tokens can be completed by speaker intentions, in which case they communicate complete generalisations of varying quantificational force and flavour via implicatures. For example, in one context, the speaker uttering (4) completes the propositional radical in such a way so as to convey that most seagulls swoop down to steal food, whereas in another, the same propositional radical can be completed by different intentions so as to convey the claim that more seagulls than other birds engage in this kind of behaviour.

4.3 Non-specificity as a by-product

Both those that posit contextualism and those that posit semantic incompleteness in their accounts of generics emphasise the contexts of utterance and, in particular, the intentions of speakers to explain the content of generics utterances. This emphasis also means that they have the resources to understand non-specific uses of generics as uses that arise out of a *certain kind of context of utterance*, namely one that lacks those intentions or other ingredients that determine one specific generalisation as what is communicated by a generic token.¹⁷ Following Kai von Fintel and Anthony Gillies (2011), who discuss related issues surrounding epistemic modals, I will call such contexts 'indeterminate contexts'.

¹⁵ Bach (2006, p. 436).

¹⁶ An implicature is content that is conveyed pragmatically, much like a Gricean implicature, but more closely related to the semantic content of the utterance. See Bach (1994).

¹⁷ Alternatively, the relevant contexts may contain conflicting ingredients. I ignore such cases for the sake of brevity, but the considerations that apply to them are similar.

In describing the type of context, Sterken says:

The speaker needn't intend that a determinate generalisation is expressed by Gen. The conversation may not demand this, rather the speaker might simply intend some range of generalisations, so long as the speaker and hearer manage to coordinate on or attend to an appropriate generalisation. (Or, an appropriate range of quantifiers that act as the quantificational force, and an appropriate range of properties that can act as the domain) (2015a, pp. 21–22).¹⁸

Both Nguyen and Sterken use the term 'indeterminate' to describe utterances of generics like Aoife's, although Sterken also describes them as 'underspecified'. I think that it is important to distinguish these. The two options on the table are that either:

- (a) The generic tokens express one candidate generalisation, but it is indeterminate, in that there is no fact of the matter about which one it is.
- (b) The generic tokens express non-specific or 'underspecified' generalisations

I believe that option (a) is implausible. First, picking any one of the candidate generalisations over the others is arbitrary. Second, if the speaker doesn't intend any one candidate generalisation to be communicated and the context lacks other ingredients to do the job then why believe that any one non-generic generalisation is actually expressed? I return to this worry later on in this paper. Instead, I want to explore option (b), according to which the content expressed by Aoife's utterance itself is non-specific, for the remainder of this paper. How can Sterken and Nguyen account for non-specific generic tokens? I will consider two options: that they convey no propositions or that they convey several.

5 Pragmatic explanations of non-specificity

5.1 No propositions

If generics rely on some features of the context in which they are uttered for their meaning, whether through context-sensitivity or semantic incompleteness, then, if those features are lacking, we might think that those tokens just don't express any propositions. Maybe generic tokens, at least those uttered in indeterminate contexts, just aren't the kind of sentences that express truth-evaluable, complete propositions? I will call this view 'propositional nihilism'.

Both contextualism and semantic incompleteness theories are compatible with propositional nihilism, given additional assumptions. I focus on what explanation the semantic incompleteness proponent could give. According to Nguyen, the propositional radicals that generics express need to be completed by speaker intentions in order to convey full propositions via implicatures. Thus, if the speakers lack those fine-grained intentions that would complete the propositional radicals, then no full

¹⁸ Nguyen (2019, p. 1315) also discusses non-specific uses of generics. He considers whether it is a counterexample to his view that it allows such 'indeterminate' uses of generics as a result of speakers not intending any one way in which the propositional radical is to be completed. Rightly, in my view, he points out that such uses of generics are common and thus that their possibility is not a point of criticism.

proposition at all is conveyed. So, in our example, if Aoife lacks the fine-grained intentions usually required to communicate overtly quantified generalisations, then no such generalisations can be conveyed.

However, this raises questions about how her utterance can be meaningful. If generic tokens themselves only express propositional radicals, any truth-apt generalisations must be conveyed via implicatures. But these implicatures require that speakers have fine-grained intentions to complete the propositional radicals in one of several ways. When they don't, it is unclear how any propositions could be conveyed via implicatures.

Tokens of other examples of supposed propositional radicals that are not completed by speaker intentions have a distinct *incomplete* feel to them. For example, if I say 'It is raining' with no intention of conveying that it is raining in any one location, then whatever meaning I did convey has a distinct, incomplete feel to it. However, utterances like Aoife's do not have such an incomplete feel.¹⁹ If I utter 'Dogs bark' without intending to make a claim that is specifically about, say, prevalence or what's characteristic for dogs, then that generalisation may be rough, but not incomplete in the sense that my utterance of 'It is raining' is.²⁰

Given the role generic tokens expressing rough generalisations play in our everyday life, it would seem strange if nothing they said (either semantically or via implicatures) was truth-evaluable and we would need a reason for why they are used in this way. Consequently, I will explore an alternative response, namely that generic tokens uttered in indeterminate contexts do express propositions, in fact, that they express several.

5.2 Several propositions

Propositional pluralism is the view that some utterances express several propositions. Applied to the case of generics, the idea would be that at least some generic tokens express several propositions, corresponding to the candidate generalisations. Emanuel Viebahn (2019) distinguishes between two types of propositional pluralism: strong and flexible pluralism. Both entail the claim that some sentences express sets containing several propositions, but they differ in whether the pluralism comes about as

¹⁹ Recanati (2004) discusses utterances of this type (actually including utterances of 'It is raining') as a reason against thinking of the relevant sentences as semantically incomplete by employing what he calls an 'optionality criterion'. His idea is that when sentences can be used to express what seem like complete propositions in indeterminate contexts, this is a sign that they are not semantically incomplete. My point here is not just that generics aren't semantically incomplete, but rather that even if they were we would need an explanation of how the truth-evaluable content would be communicated pragmatically in the absence of specific communicative intentions.

²⁰ In response, an anonymous reviewer for this journal suggests that indeterminate contexts already contain a lot of common ground about the kind and properties in question and that these background assumptions in some way fill in the meaning such that the utterances sound less incomplete. If we instead consider generic utterances about made up kinds and properties, like *Harpies flibbet*, these utterances might have more of an incomplete feel to them than the case I consider. I agree. However, the question for me is not whether extra content is communicated through the interaction of the content communicated and prior beliefs, but *how* this is done. If it is done through implicatures, then we need an account of how these are generated by non-specific communicative intentions. If they are not communicated through implicatures, but for instance, implicatures, then we come back to the idea that by itself, my utterance of 'Dogs bark' appears to say much less that is directly related to the meaning of the words used than it seems. So, this response only postpones the issue.

a result of the sentence containing context-sensitive expressions (flexible pluralism) or semantically incomplete expressions (strong pluralism). This means that pluralism is compatible with both contextualism and semantic incompleteness theories. I will be focusing on flexible pluralism here, but analogous considerations apply for strong pluralism.

According to flexible pluralism, sentences that include context-sensitive expressions express a set containing several propositions when uttered in indeterminate contexts. To my knowledge, no-one has yet proposed such an account for generics. According to flexible pluralism about generics, generic tokens uttered in indeterminate contexts express several non-generic generalisations. Generic tokens always express a set of propositions, but this set can contain more than one element when uttered in indeterminate contexts.

Pluralism does a good job of explaining non-specific uses of generics. Generic tokens uttered in indeterminate contexts don't just express one fine-grained generalisation, they actually express several. We don't need to choose between the different candidate generalisations Aoife's utterance might have expressed, because it actually expressed all contextually salient ones. Pluralism gives, at least at first sight, a straightforward account of what contents Aoife's utterance conveys in terms that posit things most already accept: sets and propositions.²¹

Propositional pluralism has its advantages, but it also comes with some costs that speak in favour of finding an alternative approach. Propositions are posited not just as semantic values of sentences, but also, for example, as the objects of propositional attitudes. To emphasise one such attitude in particular, consider generic beliefs. What we believe when we believe, or even know, that bats have good hearing, seems to be quite different from what we believe when we believe that most bats have good hearing, with regards to what evidence our beliefs require, but also what other beliefs and actions they give rise to. Does believing that bats have good hearing really require having a host of fine-grained beliefs about bats and having good hearing? What norms govern them? Are they only properly expressed as assertions if we meet epistemic requirements for all the propositions they convey? I believe that the fact that generic beliefs are non-specific can explain some of the differences between them and beliefs in other generalisations, making it all the more important that we understand the source of their non-specificity. Adopting propositional pluralism, and nihilism for that matter, comes with a large amount of extra work involved in modifying theories that rely on sentences only expressing one proposition at a time or coming up with new ones.

However, there is also another concern that I have already anticipated in my discussion of Sterken's analysis. If communicative intentions are necessary for giving context-sensitive expressions their meaning, as is the case with supplementives, *lacking* the communicative intentions required for those expressions to have a given

²¹ Another benefit of propositional pluralism is that it provides a unified account of non-specificity that can arise from sentences with different kinds of context-sensitive expressions being uttered in indeterminate contexts. For instance, Sterken (2015a) compares these uses of generics to non-specific uses of 'here' in which speakers can leave open how big an area they are referring to. Propositional pluralism can account for such similarities, simply by including 'here' and other expressions like it, in the group of expressions which can have several semantic values when used in indeterminate contexts.

meaning, should make it impossible for the expression to have that meaning.²² Given that Aoife utters (4) without having intentions to convey any one non-generic generalisation, it just seems unmotivated to think that her utterance communicates any of those candidate generalisations, much less several of them. These considerations illustrate that while propositional pluralism has the resources to account for non-specific uses of generics, it also brings with it complications. In the last section, I therefore want to propose an alternative that avoids such complications: the idea that generics express *one* proposition at a time, just a non-specific one.

6 The generality account

6.1 The basic idea

My proposal is this: rather than assume that generics generally express fine-grained generalisations, they, by default, express coarse-grained generalisations. Generics generalise about members of a kind and attribute a link between them and some property, but they need not say much about the nature of that link. In contrast to the explanations I considered in the previous section, the generality account understands non-specificity not as a by-product of context-sensitivity or semantic incompleteness, but instead simply as a feature of the semantic content of generics.²³

Though the idea that generics express non-specific generalisations hasn't been widely discussed in the generics literature, the idea behind it is actually somewhat obvious.²⁴ Generics lack overt quantifier expressions. If quantifier expressions in sentences that express generalisations determine which such generalisations they express,

²² In response, an anonymous reviewer for this journal suggests that the relevant contexts should be thought of not as lacking the relevant ingredients, but instead as containing them by default. The idea would then be that several fine-grained generalisations are communicated as long as they are not in conflict with the speaker's intention. One worry about this strategy would be that it would entail that one would often say more, semantically or via implicatures, by speaking generally than if one was more specific.

²³ An anonymous reviewer for this journal suggests that contextualists can argue that generics uttered in indeterminate contexts do express weak propositions as a result of existential closure, without needing to posit non-specificity separately. Existential closure is a process whereby variables that would otherwise be free are existentially bound (Heim, 1982). My best guess of what such an account would look like is as follows: Where generics are uttered in indeterminate contexts, whatever values Gen is compatible with in the context would be existentially quantified over. The main difference between this account and the one I describe on this page is that it would be a process that is not triggered generally, but only when generics are uttered in indeterminate contexts. I do actually agree that existential closure might be what gives rise to non-specificity, albeit triggered by an absence of an overt quantifier expression at the surface level, not the context of utterance. Two important questions that bear on this matter are: First, for this proposal to not be ad-hoc, existential closure is something that should occur with any expressions that are context-sensitive/semantically incomplete when they are uttered in indeterminate contexts. The proposal for generics would depend on whether this is a plausible general strategy. Second, I will suggest that context-sensitivity may not be needed to account for token variability of the kind that motivates contextualists in the next few pages, in which case we might take non-specificity as a theoretical starting point and do without contextualism.

²⁴ The most similar other accounts I am aware of are by Nickel (2016) who argues that generics existentially quantify over ways of being normal, Carlson (2008) who argues that generics quantify existentially over inductive patterns, Greenberg (2004, 2007) who argues that they are vague, and Lemeire (2020) who argues that they express disjunctive generalisations.

then it doesn't seem far-fetched to think that sentences that lack quantifier expressions express less restricted generalisations. So how do generics encode non-specificity semantically?

According to the generality account, generics encode non-specificity by existentially quantifying over non-generic generalisations concerning the kind and property in question.²⁵ So, Aoife's utterance existentially quantifies over non-generic generalisations concerning seagulls and the swooping down to steal food: that most seagulls do it, that more seagulls than other types of bird do it, that it is normal for them to do so, and so on. Consequently, (4) is true as long as *at least one* non-generic generalisation about seagulls swooping down to steal food is true.

Generic generalisations still generalise about members of a kind by linking it to some property but are less specific about the nature of that link than non-generic generalisations. Not just any kind of link between the kind and property in question will do, however. For example, if most seagulls were scared of swooping down to steal food and none ever in any sense did it, this would not suffice to make (4) true. Instead, the domain of generalisations about seagulls that are eligible to make true the generic needs to be restricted.

6.2 The domain of quantification

Exactly which non-generic generalisations generics quantify over is not a matter I can or want to stipulate at this point. Doing so would require empirical work that hasn't yet, as far as I am aware, been undertaken. What is needed is experimental testing of which non-generic generalisations speakers take to make true the corresponding generics. These seem likely to me to include those generalisations in terms of which generics have been analysed: claims about what's normal for members of a kind, probability claims (both absolute and comparative), claims about what dispositions members of a kind have, claims about the essence of members of a kind, and perhaps claims about what norms members of a kind are subject to. However, in the absence of adequate data, it would be premature to limit which non-generic generalisations can feature in the domain of quantification.

However, some, to my mind plausible hypotheses for restriction have been given by Prasada and Dillingham (2006) and Nickel (2016), who emphasise the importance of *explanation* in what can make true generics. For example, we might think that for (7) to be made true, it needs to be the case that being a horse in some way explains the wearing of horseshoes. Nickel (2016) in particular, emphasises that different types of links can be explanatory depending on the type of inquiry in which one is engaged, which would explain the variety in non-generic generalisations that can make true a generic.

One thing that is important to point out is that each of the generalisations generics quantify over existentially can themselves be non-specific to varying degrees. Some

²⁵ While I talk about generics as quantifying existentially over a domain of non-generic generalisations for the remainder of this article, I remain neutral here as to how this quantificational force arises. See Sterken (2016) for arguments in favour of quantificational approaches that posit a covert quantifier and Liebesman (2001), Cohen (2012), and Collins (2018) for alternative, pragmatic explanations.

generalisations contain a lot of information about the nature of the link they ascribe, some very little, and most are somewhere in between. What is important for my purposes is that the generalisations in question must be *more* fine-grained than generics in that they specify more of the flavour or quantificational force of the link between the kind and property in question.

- (8) 87% of Scots are vitamin D deficient
- (9) Most Scots are vitamin D deficient
- (10) Many Scots are vitamin D deficient

For example, (8) is a generalisation about Scottish people and vitamin D deficiency. The generalisation is more specific than (9), which in turn is more specific than (10). The less specific the generalisation, the less information is specified about the nature of the link it posits between a kind and a property. The idea now is that generics are at the non-specific end of the specificity scale for generalisations: they are the non-specific counterparts to other generalisations.

The generality account can explain non-specificity in generics without needing to pose phenomena like context-sensitivity or semantic incompleteness simply as a feature of their semantic content. The view also has other advantages: it can provide the resources for a straightforward picture of generic beliefs, explain the varying strength of negated generics in terms of scope ambiguities, and may be able to explain why generic sentences are, as sentences featuring other non-specific expressions, associated with particular types of tense, aspect, and mood.²⁶ Much else needs to be said to develop this approach and expand on these advantages, however, I want to end by responding to what I expect will be the most pressing objection to it: doesn't the generality account overgenerate true generics?

7 Overgeneration

According to the generality account, as long as one of a set containing several non-generic generalisations about the kind and property in question is true, so is the generic. However, there are many cases in which a non-generic generalisation about a kind and property is true, but the corresponding generic seems false.

For example, consider the widely discussed case of:

- (11) Books are paperbacks

Even though most books are paperbacks, utterances of (11) still sound strange and, to some, simply false.²⁷ However, according to the generality account, (11) should be true because most books are paperbacks. The problem isn't just limited to sentences like (11), but also arises in a more limited form about generic tokens in particular contexts of utterance. What led Sterken and Nguyen to posit context-sensitivity or semantic incompleteness was the observation that tokens of generics appear to express *different* generalisations across contexts of utterance. But in its invariant version, the

²⁶ See Frawley (2013, pp. 69–74) on grammatical markers of non-specificity and Dahl (1995) and Chierchia (1995) on such markers in generics.

²⁷ For example, Leslie (2008).

generality account predicts that the truth-conditions of generics are stable because they existentially quantify over the same set of non-generic generalisations.

7.1 Contextualism to the rescue?

One straightforward way of warding off both these worries is to combine the generality account and contextualism. Instead of arguing that generics existentially quantify over the *same* set of non-generic generalisations across all contexts of utterance, they might instead quantify over a *contextually restricted* set of them. Such an account could explain token variability in the same way other contextualist accounts do, while giving a straightforward answer to the question about the propositional content of non-specific generic tokens.

The contextualist picture would be as follows: generic sentence types are highly non-specific. They simply say that at least one contextually salient non-generic generalisation about members of the kind and the property in question is true, but little more about the nature of that generalisation. When generic types are tokened, their level of specificity depends on the contexts in which they are uttered. The less restrictive a context is, the more flavours and strengths of association are salient and, hence, the less specific the content of the generic token.

The contextualist version of the generality account has a response to overgeneration worries. Many of the generic tokens that the invariant version of the generality account would controversially predict to be true, could be false according to the contextualist version simply because the non-generic generalisation that is true would not be contextually salient. It can also account for token variability in the standard way contextualists do. A generic can be used to convey a generalisation of a particular flavour or quantificational force in one context. Yet in another context, the same generic type could be used to express a generalisation of a different flavour or force because the domain of quantification would contain different non-generic generalisations for the generic to quantify over.

7.2 Embracing true generics

The contextualist view would alleviate some of the urgency of overgeneration worries, at least initially. However, I don't think that this is a conclusive point in its favour for three reasons: the relief it provides is temporary, it is not clear that accepting many generics as true is something we should avoid, and accepting them as true doesn't preclude us from explaining why utterances of them can sound bad.

First, the overgeneration problem simply resurfaces for the contextualist as a metasemantic question about why a given non-generic generalisation cannot be salient in a given context. Why can't we utter (11) intending to make the claim that most books are paperbacks and thereby convey it?

Second, the existing evidence we have about the intuitions speakers have about many of the generics assumed to be straightforwardly false, is quite limited. The studies we do have show a complicated picture with regards to people's truth-value

judgements.²⁸ Even within the generics literature, there is disagreement about the truth-conditions of some generics. For example, Krifka et al. (1995) assume that it is important to make sure a theory of generics doesn't predict that *Children born in Rainbow Lake are right-handed* is true. However, Wasserman (2011) casts doubt on this assumption. Likewise, Leslie (2007, 2008) assumes that *Sharks attack bathers* is true, whereas Sterken (2015b) argues that it, and a class of other generics Leslie takes to be true, are actually false. Specifically, with regards to (11), *Books are paperbacks*, Leslie (2008) assumes that it is false, whereas Nickel (2018) argues that it is not.

For several reasons then, more evidence is required to substantiate the idea that the relevant generics are actually false before being concerned that the generality account or other non-specificity approaches, overgenerate true generics. These considerations may do something to lessen the force of worries about overgeneration with regards to sentences like (11). But what about the ability to explain the observations of what looks like variability in the truth-conditions of generic tokens?

In the last section of this paper, I will argue that we can explain the ways in which generics can be *used* to convey a variety of more specific generalisations, even if their truth-conditions are stable and non-specific. Even though generic tokens *semantically* express non-specific generalisations, they, in addition, convey more specific generalisations through pragmatic means. So, the examples Sterken and Nguyen cite are not instances of truth-conditional variability, but of variability in use made possible by the pragmatic effects of generic tokens.

8 Pragmatic effects

Non-specific language can convey more specific contents by means of its pragmatic effects. We can use generics to semantically express non-specific generalisations which interact with cognitive biases, prior beliefs, and conversational expectations to cause beliefs in more specific generalisations. Going back to the seagull example, even if (4) doesn't semantically tell us which non-generic generalisation(s) make(s) it true, this doesn't preclude non-semantic factors from causing me to form just such beliefs. Depending on my background assumptions about seagulls, beliefs about the topic of conversation, or cognitive biases, I might well form additional beliefs about which *proper subset* of generalisations, or even which individual generalisations are responsible for making (4) true on the basis of Aoife's utterance.

To develop this line of argument, we can make use of an existing literature on the implicatures generic utterances generate. Sally Haslanger (2011, 2014) argues that most generic tokens produce distinctive conversational implicatures.²⁹ She argues

²⁸ For example, a detailed study by Prasada et al. (2013), which asked participants to judge the truth-values of generic sentences, shows quite a mixed picture. Most of the supposedly false generics don't receive strong falsity scores. For example, (11), *Books are paperbacks*, received a mean score of 0.21 on a scale of -3 to $+3$, where -3 was definitely false and $+3$ definitely true. Even those generics that were judged as more false than true didn't receive strong scores. For example, the mean scores for the two types of supposedly false generics (those in which the majority of members have the property and those in which only a minority do) were 0.21 and -0.52 respectively.

²⁹ For Haslanger (2014, p. 15), utterances of what she calls 'statistical generics' do not generate the relevant implicatures. Because I think that generics are non-specific, I think the lack of relevant implicatures for

that cognitive biases lead us to expect that generalisations about members of a kind and a property are made true by essences shared by members of the kind, which in turn give rise to expectations that they should exhibit the property. For instance, hearing that seagulls swoop down to steal people's food may lead us to believe that seagulls have some underlying essence that grounds other, e.g., prevalence generalisations about them, and even gives rise to normative expectations about seagulls, such as that a proper seagull ought to behave in this manner.³⁰

Other important implicatures generated by generic tokens depend less on general normative or essentialist biases, and more on prior beliefs about the relevant kinds and properties, or general conversational relevance expectations.³¹ First, prior beliefs about which non-generic generalisations are *not* true can restrict the domain of generalisations that interlocutors think of as likely contenders for making true a generic. For example, most participants in conversations in which (7), *Horses wear horseshoes*, is uttered know that horses don't naturally wear horseshoes which restricts the set of non-generic generalisations that they think could make true (7). Second, when a context is such that only a proper subset of non-generic generalisations would be relevant, uttering a generic implicates that it is one of those non-generic generalisations in the restricted set that makes the generic token true.

Because of this interaction between relevance expectations and prior beliefs, speakers can use generics to pragmatically convey more specific generalisations. Even though generics don't, according to the generality account, have a strong semantic meaning, they can nevertheless pragmatically communicate that a given generic is true in virtue of a restricted subset of generalisations about a kind and property. This account of the pragmatic effects of generic utterances can explain both the appearance of token variability and why some generics the generality account predicts are true may nevertheless sound bad to us.

For example, in a conversation between horse breeders about horse-care related budgets, an utterance of (7), *Horses wear horseshoes*, may generate implicatures about the prevalence of horseshoe-wearing among horses or perhaps a normative requirement for them to do so, but likely none about how horses have evolved. In contrast, in a context in which two evolutionary biologists are discussing the various adaptations horses have undertaken in response to their environment, uttering (7) would likely generate the implicature that horses naturally wear horseshoes. Such an utterance would sound strange, not because it is false, but because the generalisation that does make it true is not of the kind that is relevant in the context of utterance and the generalisation that would be relevant is false.

Footnote 29 continued

some generics needs to be accounted for by extra-semantic factors, for example, additional beliefs about the kind and property in question.

³⁰ Haslanger's argument connects to a literature by Susan Gelman and various co-authors who have drawn attention to what they call the essentialising effects of generics, i.e., the tendency of generics to lead us to expect that more superficial properties are grounded in essences shared by members of a kind. See Gelman et al (2010), Leslie (2014), Cimpian et al. (2010), Gelman (2003). This phenomenon might also be a feature of non-generic generalisations. See Hoicka et al. (2018) and Munton (2019).

³¹ For related work on generic implicatures, see Saul (2017) who describes how expectations concerning relevance can make generic tokens sound infelicitous when the link they posit is irrelevant as well as Rosola (2019) who argues that generics convey particularised conversational implicatures depending on the investigative aims of a context.

Generics provide particularly fertile grounds for implicatures because of their weak semantic contents. By positing a connection between kinds and properties, they provide enough contents to interact with cognitive biases, prior beliefs, and conversational relevance expectations. At the same time, what they semantically express is non-specific enough to be compatible with many possible implicatures and thereby allows them to be generated. In this way, the generality account can explain the semantic basis that facilitates pragmatic effects in generic tokens.

There is much further work to be done in identifying the underlying cognitive and pragmatic mechanisms that generate generic implicatures and the ways they interact. However, at least for now, it seems that there are pragmatic resources available to the proponent of the generality account that lets them address overgeneration worries and explain use variability without requiring contextualism.

9 Conclusion

I have drawn attention to what I take to be an underappreciated aspect of generics: their non-specificity. I considered views according to which non-specificity occurs as a by-product of contextualism or semantic incompleteness but argued that these approaches commit us to propositional nihilism or pluralism, both of which are undesirable and unnecessary. Instead, I proposed a novel account, according to which generics quantify existentially over non-generic generalisations. According to the generality account, non-specificity is a feature of the semantic content of generic sentences. A generic is true as long as one of several non-generic generalisations about the kind and property in question is. I've provided an initial defence of the account against the worry that it overgenerates true generics and argued that though generics semantically express weak generalisations, their pragmatic effects are anything but.

Acknowledgements I'd like to thank Rachel Fraser, Richard Holton, Paula Keller, James Kirkpatrick, Rae Langton, Lukas Lewerentz, Benjamin Marschall, Lucy McDonald, Martina Rosola, Rachel Sterken, Emanuel Viebahn as well as anonymous reviewers for this journal for discussion and comments. Special thanks to Aoife Deery for providing the seagull example. I presented ancestors and other close relatives of this paper at the *Just Words* workshop at University College London, the *Metasemantics of Context-sensitivity and Ambiguity workshop* at Humboldt-Universität Berlin, The *Oxford Graduate Conference*, and The *Indeterminacy and Underdetermination Conference* at University College Dublin and would like to thank participants at these events for their questions. My research was supported by a studentship from the Arts and Humanities Research Council.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anderson, L., Haslanger, S., & Langton, R. (2012). Language and Race. In *Routledge Companion to Philosophy of Language*. Routledge.
- Asher, N., & Morreau, M. (1995). What some generic sentences mean. In G. Carlson & F. J. Pelletier (Eds.), *The generic book* (pp. 300–338). University of Chicago Press.
- Asher, N., & Pelletier, F. J. (2012). More truths about generic truth. In *Genericity* (pp. 313–333). Oxford University Press.
- Bach, K. (1994). Conversational implicature. *Mind & Language*, 9(2), 124–162.
- Bach, K. (2006). The excluded middle: Semantic minimalism without minimal propositions. *Philosophy and Phenomenological Research*, 73(2), 435–442.
- Carlson, G. (2008). Patterns in the semantics of generic sentences. In J. Guéron & J. Lecarme (Eds.), *Time and modality* (pp. 17–38). Springer.
- Chierchia, G. (1995). Individual level predicates as inherent generics. In G. Carlson & F. J. Pelletier (Eds.), *The generic book* (pp. 176–223). Chicago University Press.
- 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.
- Cohen, A. (1999a). *Think generic! The meaning and use of generic sentences*. Carnegie Mellon University.
- Cohen, A. (1999b). Generics, frequency adverbs, and probability. *Linguistics and Philosophy*, 22(3), 221–253.
- Cohen, A. (2012). No quantification without reinterpretation. In A. Mari, C. Beyssade, & F. Del Prete (Eds.), *Genericity*. Oxford University Press.
- Collins, J. (2018). Genericity sans gen. *Mind & Language*, 33(February), 34–64.
- Dahl, Ö. (1975). On generics. In E. L. Keenan (Ed.), *Formal semantics of natural language* (pp. 99–111). Cambridge University Press.
- Dahl, Ö. (1995). The marking of the episodic/generic distinction in tense-aspect systems. In G. Carlson & F. J. Pelletier (Eds.), *The generic book* (pp. 412–425). Chicago University Press.
- Delgrande, J. P. (1987). A first-order conditional logic for prototypical properties. *Artificial Intelligence*, 33(1), 105–130.
- Eckardt, R. (2000). Normal objects, normal worlds and the meaning of generic sentences. *Journal of Semantics*, 16(3), 237–278.
- Frawley, W. (2013). *Linguistic semantics*. Routledge.
- Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. Oxford University Press.
- Gelman, S. A., Ware, E. A., & Kleinberg, F. (2010). Effects of generic language on category content and structure. *Cognitive Psychology*, 61(3), 273–301.
- Greenberg, Y. (2004). *Manifestations of genericity*. Taylor & Francis.
- Greenberg, Y. (2007). Exceptions to generics: Where vagueness, context dependence and modality interact. *Journal of Semantics*, 24(2), 131–167.
- Haslanger, S. (2011). Ideology, generics, and common ground. In C. Witt (Ed.), *Feminist metaphysics: Explorations in the ontology of sex, gender and the self* (pp. 179–207). Springer.
- Haslanger, S. (2014). The normal, the natural and the good: Generics and ideology. *Politica & Società*, 36(5), 392.
- Heim, I. (1982). *The semantics of definite and indefinite noun phrases*. UMass Amherst.
- Hoicka, E., Saul, J., Prouten, E., Whitehead, L., & Sterken, R. (2018). Language signalling high proportions, not just generics, leads to essentializing for novel social kinds. Psyarxiv.com. <https://doi.org/10.31234/osf.io/xe6sj>.
- King, J. C. (2013). Supplementives, the coordination account, and conflicting intentions. *Philosophical Perspectives. A Supplement to Nous*, 27(1), 288–311.
- Krifka, M., Pelletier, F. J., Carlson, G., TerMeulen, A., Chierchia, G., & Link, G. (1995). Genericity: An introduction. In G. Carlson & F. J. Pelletier (Eds.), *The generic book* (pp. 1–124). University of Chicago Press.
- Lemeire, O. (2020). Falsifying generic stereotypes. *Philosophical Studies*, 178(7), 2293–2312.
- Leslie, S.-J. (2007). Generics and the structure of the mind. *Philosophical Perspectives. A Supplement to Nous*, 21(1), 375–403.
- Leslie, S.-J. (2008). Generics: Cognition and acquisition. *The Philosophical Review*, 117(1), 1–47.

- Leslie, S.-J. (2012). Generics articulate default generalizations. *Recherches Linguistiques De Vincennes*, 41, 25–44.
- Leslie, S.-J. (2014). Carving up the social world with generics*. In *Oxford studies in experimental philosophy* (pp. 208–231). Oxford University Press.
- Leslie, S.-J. (2015). 'Hillary clinton is the only man in the obama administration': Dual character concepts, generics, and gender. *Analytic Philosophy*, 56(2), 111–141.
- Leslie, S.-J. (2017). The original sin of cognition: Fear, prejudice, and generalization. *The Journal of Philosophy*, 114(8), 393–421.
- Liebman, D. (2011). Simple generics. *Noûs*, 45(3), 409–442.
- McConnell-Ginet, S. (2012). Generic predicates and interest-relativity. *Canadian Journal of Linguistics/revue Canadienne De Linguistique*, 57(2), 261–287.
- Munton, J. (2019). Beyond accuracy: Epistemic flaws with statistical generalizations. *Philosophical Issues. A Supplement to Nous*, 29(1), 228–240.
- Nguyen, A. (2019). The radical account of bare plural generics. *Philosophical Studies*, 177(5), 1303–1331.
- Nickel, B. (2008). Generics and the ways of normality. *Linguistics and Philosophy*, 31(6), 629–648.
- Nickel, B. (2016). *Between logic and the world: An integrated theory of generics*. Oxford University Press.
- Nickel, B. (2018). Generics, conservativity, and kind-subordination. *Philosophers' Imprint*, 18(13), 18.
- Prasada, S. & Dillingham, E. (2006). Principled and statistical connections in common sense conception. *Cognition* 99(1), 73–112.
- Prasada, S., Khemlani, S., Leslie, S.-J. & Glucksberg, S. (2013). Conceptual distinctions amongst generics. *Cognition* 126(3), 405–22.
- Recanati, F. (2004). *Literal meaning*. Cambridge University Press.
- Ritchie, K. (2019). Should we use racial and gender generics? *Thought A Journal of Philosophy*, 8(1), 33–41.
- Rosola, M. (2019). Generics in context: The robustness and the explanatory implicatures. In G. Bella & P. Bouquet (Eds.), *Modeling and using context* (Vol. 11939, pp. 223–237). Springer International Publishing.
- Saul, J. (2017). Are generics especially pernicious? *Inquiry*, February 1–18.
- Sterken, R. K. (2015a). Generics in context. *Philosophers' Imprint*, 15(21), 30.
- Sterken, R. K. (2015b). Leslie on generics. *Philosophical Studies*, 172(9), 2493–2512.
- Sterken, R. K. (2016). Generics, covert structure and logical form. *Mind & Language*, 31(5), 503–529.
- Viebahn, E. (2019). *Semantic Pluralism*. Klostermann.
- von Fintel, K., & Gillies, A. (2011). Might made right. In A. Egan & B. Weatherson (Eds.), *Epistemic modality*. Oxford Scholarship Online.
- Wasserman, R. (2011). Dispositions and generics. *Philosophical Perspectives. A Supplement to Nous*, 25(1), 425–453.