Generics in Use

Anne Bosse

Peterhouse

University of Cambridge

This thesis is submitted for the degree of Doctor of Philosophy

July 2021

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text.

It is not substantially the same as any work that has already been submitted before for any degree or other qualification except as declared in the preface and specified in the text. I further state that no substantial part of my dissertation has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text.

It does not exceed the prescribed word limit for the relevant Degree Committee.

Abstract

Generics in Use

Anne Bosse

This thesis is about generics, sentences like 'Bricks are red', 'Boars have bristly hair', or 'British people love peas'.

In Chapter 1, I consider what makes such sentences generics. I propose that generics should be defined as generalisations that lack overt quantifier expressions.

In Chapter 2, I question an assumption made in much of the generics literature, namely that generics express specific generalisations. I consider explanations according to which non-specificity in generics is a by-product of context-sensitivity or semantic incompleteness, but instead propose that generics semantically express non-specific generalisations by default.

In Chapter 3, I present a novel account of the semantics of generics that can explain their non-specificity. According to the generality account, generics of the form *Fs are G* are true iff at least one of several non-generic generalisations about the kind F and the property G is true.

In Chapter 4, I turn to the mental states generics give voice to. I evaluate Sarah-Jane Leslie's account, according to which generics express cognitively basic generalisations and propose an alternative. Just as we can express non-specific generalisations in speech using generic sentences, we can also take various propositional attitudes, including of belief, towards them.

In Chapter 5, I consider what functional role genericity plays in thought and speech. I argue that because generic beliefs are non-specific, they are also evidentially undemanding: they take whatever evidence they can get. This allows us to form inferentially powerful beliefs and even gain knowledge about kinds based on limited evidence.

In Chapter 6, I focus on the connection between generics and stereotyping. I propose an account of stereotypes according to which they involve generic beliefs. I end by considering how this analysis can inform our responses to stereotyping.

Contents

Declaration 3

Abstract 5

Acknowledgements 9

Introduction 11

Chapter 1: What are Generics? 15

1 Introduction 15
2 Defining generics: two options 16
3 The problem 19
4 Defining generics as unmarked 22
5 Quantification in generics pt. 1: The case against 25
6 Quantification in generics pt. 2: The case for 27
7 Unmarkedness and covert quantifiers 30
8 Conclusion 32

Chapter 2: Non-specific generics 35

1 Introduction 35
2 Truthmaker variability in the wider literature 36
3 The specificity assumption and a counterexample 38
4 Token variability 40
5 Explaining token variability 42
6 Propositional nihilism 45
7 Propositional pluralism 47
8 Looking for an alternative view 48
9 Conclusion 49

Chapter 3: The Generality Account 51

1 Introduction 51
2 An alternative proposal 51
3 Existential closure 54
4 Comparison with Nickel 56
5 Advantages 60
6 Overgeneration 62
7 Pragmatic contents 66
8 Conclusion 69

Chapter 4: Generic Belief 71

1 Introduction 71
2 Generics as default hypothesis 71
3 Issues with Leslie's account 75
4 Generic belief 80
5 Generic inferences 84
6 Conclusion 87

Chapter 5: What generics are for **89**
1 Introduction 89
2 Functions 89
3 Generic beliefs as evidential omnivores 92
4 Making the most of what evidence we have 94
5 Conceptual role 97
6 In speech 100
7 Conclusion 103

Chapter 6: Generics and Stereotyping **105**
1 Introduction 105
2 What is stereotyping? 105
3 What's the connection between generics and stereotyping? 107
4 The proposal: stereotypes involve generic beliefs 108
5 The inner workings of stereotyping 109
6 So, is it a stereotype because it's true? 113
7 A wrong in stereotyping 115
8 Conclusion 119

Conclusion **121**
Looking back 121
Further work 122

Bibliography **125**

Acknowledgements

The research that led to this thesis was made possible by generous funding from the Arts and Humanities Research Council.

My first thanks must go to my supervisors, Rae Langton and Richard Holton. Rae, it has been genuinely inspiring to have had you as my supervisor. You have taught me a lot about how philosophy should be done. Richard, thank you so much for all your insight, open-mindedness, and encouragement. This thesis would've been much worse without you. I would also like to thank support staff in the Faculty of Philosophy, in particular, Jo Harcus and Charlie Evans for helping with any admin or library-related queries, and for creating a friendly atmosphere.

There are many academics from whose insight this thesis has benefitted: Rachel Fraser, Alex Jackson, James Kirkpatrick, Lukas Lewerentz, Sally McConnell-Ginet, Jessie Munton, Anthony Nguyen, Martina Rosola, Jenny Saul, Robert Simpson, Paulina Sliwa, Rachel Sterken, Emanuel Viebahn, and Alnica Visser. I also would like to thank anonymous reviewers for the Australasian Journal of Philosophy and Synthese for helpful comments on papers based on Chapters 2 and 6.

I have presented what were to become chapters of this thesis at different venues. Though I cannot thank every helpful participant, or even note each event individually, I'd like to mention the Words (formerly Meaning) workshop, the Oxford Graduate Conference, the Generic workshop, and the Moral Sciences Club at Cambridge in particular.

I owe a huge debt of gratitude to those (then) at the Universities of Glasgow, Stirling, and St Andrews without whose influence and encouragement I probably would not have pursued a PhD in the first place. Particular thanks are due to Fraser Macbride and Katherine Hawley for their support and mentorship.

I've had an excellent graduate community in Cambridge. Thanks to Senthuran Bhuvanendra, Daisy Dixon, Matt Dougherty, Matilda Gillis, Cathy Mason*, Li Li Tan, and Zoe Walker for some insight and a lot of friendship. Particular thanks go to Paula Keller, Benjamin Marschall, and Lucy McDonald who read countless drafts of my work and provided companionship and support throughout my PhD.

I would also like to thank my family: Mama und Papa, ich möchte mich bei euch für eure Unterstützung und Offenheit bedanken. In eurer Rechtschaffenheit, Umsicht, Uneigennützigkeit und Achtung anderer seid ihr immer mein Vorbild. Kai, Jo, Omi und Geli, vielen lieben Dank für eure stetige Unterstützung und euer Interesse. Susan and Wilf, thank you for welcoming me to Scotland and into your family - I'll always treasure having been a part of it.

---

* Huge thanks to Cathy Mason, Louis Strachan, and Alex Assaly for help with proofreading.

Special thanks to Alex Assaly* for his love and support throughout the second half of my PhD. I don't think you realise just how crucial all your help has been and how appreciated it is.

My final thanks are to old friends: Paul Butterfield, Aoife Deery, Claudia Strate, Louis Strachan*, and Miriam Wilson. I don't know what to say other than that I honestly cannot believe how lucky I am.

Introduction

I was led to the topic of generics by a comment my brother made during Christmas 2016. In criticising a particular purchase his then-girlfriend had made, my brother, who was working in retail at the time, said something like the following: 'Women buy anything as long as you compliment their hair'. After pointing out various women that seemed immune to this strategy and him responding that my examples weren't actually incompatible with what he had said, I found myself unsure about what kinds of evidence would present genuine counterexamples to his claim. This exchange convinced me to look further into the matter, at which point it turned out that the problem wasn't just a lack of imagination on my part and that there were genuine questions about when generalisations like the one my brother had made were true or false. This thesis is about those generalisations.

(1) British houses are damp

(2) Nurses are caring

(3) Okapis are solitary

(1)-(3), and other sentences like them, are called 'generics'. In this thesis, I develop a new account of what generics mean and what the beliefs they give voice to are like. I then argue that this proposal can provide an attractive account of the functional roles generic generalisations play in thought and speech.

Chapter 1 is about what *makes* sentences like (1)-(3) generics. I propose that generics should be defined as generalisation sentences that lack *overt* quantifier expressions. This criterion invites thorny cross-linguistic questions and suggests a picture on which genericity comes in degrees. I argue that this is a good thing as it can help us explore differences but also acknowledge continuities between different types of sentences that express generalisations.

The literature on generics as a distinctive class of sentences is relatively young. It started in the 1970s but has taken off quickly since the publication of *The Generic Book* in 1995.[2] That book set out a research project, namely, to formulate a semantic theory for generics. Though a lot of progress has been made, those engaged in the project still struggle with many of the same problems identified in the introduction of the book. I discuss one central challenge, namely to account, by way of a semantic theory, for the large variety of ways in which generics can be made true, in Chapter 2.

Generic sentences can be made true by many different kinds of facts: facts about what most members of a kind are like, facts about what they are supposed to be like, facts about what is characteristic of them, and even facts about what norms they are subject to. For example, (1) might be true in virtue of the fact that more British than Continental European houses are damp, (2) in virtue of the fact that nurses should be caring, and (3) in virtue of the fact that being solitary is characteristic of okapis. Further, utterances of

[2] To my knowledge, starting with Milsark (1974) and accelerating after work by Carlson (1977).

one and the same generic sentence type can appear to be true in virtue of different types of facts across different contexts of utterance. For example, an utterance of (2) can seem to be true in virtue of the fact that nurses should be caring in one context, but another utterance of (2) can seem to be true in virtue of the fact that most nurses are caring. This flexibility has made it difficult to give uniform theories about the truth-conditions of generics and, consequently, led some to posit phenomena like context-sensitivity or semantic incompleteness.

In this thesis, I give a novel explanation of this flexibility that eschews those strategies. I argue that generics have an under-explored feature, namely that they are non-specific. Utterances of (2) that are true in virtue of *either* the fact that most nurses are caring, that they should be caring, or some other generalisation, are the norm, not the exception. The truth-conditions of generics have been thought elusive because those looking for them have been conducting their search at too high a level of specificity. Generics are just not the types of sentences that express generalisation that are specific with regards to their quantificational force and flavour.

With the groundwork set in Chapters 1 and 2, I develop my own, novel account of generics which can explain their non-specificity in Chapter 3. According to the account I propose, the 'generality account', a generic is true if and only if at least one of several non-generic generalisations about members of the kind and the property in question is true but does not further specify the nature of that generalisation. Generics do this by *existentially quantifying over non-generic generalisations*. For example, (2) just says that at least one of a set containing several non-generic generalisations about nurses and being caring is true but doesn't say anything more about which generalisations that is.

Despite semantically expressing highly weak generalisations, I argue that generic utterances can pragmatically convey more specific generalisations by exploiting various cognitive biases and their interaction with prior beliefs and relevance expectations. So, while a generic utterance only says that at least one of several non-generic generalisations about a kind and a property are true, it nevertheless can dispose hearers to additional beliefs about *which* non-generic generalisation makes true the generic in question. This can explain how generics can be used to convey more specific generalisations while also preserving stable semantic contents that can be the object of distinctive propositional attitudes. In the second part of the thesis, I illustrate some explanatory benefits of the generality account.

Chapter 4 is about generic beliefs. I evaluate Sarah-Jane Leslie's account, according to which generics express cognitively primitive, default generalisations. I argue that the empirical data Leslie relies on to motivate her account is much less conclusive than she presents it as being. Further, her account posits a disconnect between genericity in cognition and communication that limits its explanatory value. I suggest an extension of the generality account as an alternative. Just as non-specific generalisations can be the content of generic sentences, they can also be the objects of various propositional attitudes, including belief. On the account I propose, generic beliefs are simply attitudes towards the *same non-specific generalisations* expressed by generic sentences. This account of generic beliefs affords a better

understanding of some of the ways in which genericity is useful in thought and language, the subject of the last two chapters.

In Chapter 5, I explore how generics enable us to meet particular epistemic and communicative needs. I argue that generic generalisations enable speakers to believe, know, and assert generalisations on the basis of *limited evidence*. Because generic beliefs are non-specific, they are easily made true and can be supported by many different types of evidence. This is useful as the evidence we have access to often would not be enough to justify beliefs in non-generic generalisations but suffices to justify generic beliefs. As the semantic contents of generics are weak, speakers can comply with epistemic requirements for asserting them when they lack evidence required to properly assert overtly quantified generalisations, making generic sentences a type of hedging speech in which the content, not the force, of an utterance is hedged.

My aim in discussing these functions of genericity is two-fold: first, to constrain theories of genericity and second, to illuminate our understanding of various phenomena in which it plays a role. In particular, I hope that the discussion provides abductive support for the generality account insofar as it offers a fruitful explanation of some of the ways in which genericity features in thought and speech.

I apply this approach to one area that has garnered a lot of attention in Chapter 6: the connection between generics and stereotyping. I present a novel account of stereotypes as involving generic beliefs. I respond to a concern about my proposal, namely that it predicts that many stereotypes are true. I argue that what's troubling about stereotypes is not that they are false, but that they lead us to additional, harmful beliefs, both in non-generic generalisations and about individual members of social groups. I end by considering what the approach I have proposed can tell us about how we should respond to instances of stereotyping.

Chapter 1: What are Generics?

1 Introduction

Here are some examples of generic sentences, also called 'generics':

(1) Nurses are caring

(2) A heatwave causes water shortages

(3) The child learns to read before it learns to write

Is there one thing that (1)-(3) have in common which *makes them* generic sentences? If so, what is it?

Although the literature on generics is young, a large research programme on them has taken off since the publication of *The Generic Book* in 1995, which has only gathered speed during the last decade. However, as Greg Carlson points out, one issue in the study of generics 'which has over the years received relatively spotty attention, is that of what counts as a "generic"' (2008:20).

Work on generics often introduces the notion of a generic by using examples, just as I have done in this chapter.[3] This strategy avoids the definitional question in its entirety; however, it also presumes that we have some prior grasp of what generics are that can be invoked by examples. This is problematic for two reasons.

First, some types of sentences that used to be used as examples of generics no longer are. For instance, in their seminal article on generics that shaped the research programme as it exists today, Manfred Krifka et al (1995) mention the following variety of sentence types as sentences that can be generics:

> '[sentences containing] proper names, definite singular NPs [noun phrases], indefinite singular NPs, quantified NPs, bare plural NPs, and bare singular NPs […][4]
>
> Some examples of such sentences given are:
>
> (19) a. John / My brother drinks whiskey.
>
> b. A professor drinks whiskey.
>
> c. Every professor drinks whiskey.
>
> d. Professors drink whiskey.
>
> e. Milk is healthy.' (1995:8)

[3] For example, see Haslanger (2011, 2014), Leslie (2008, 2017), Liebesman (2011), Nickel (2018).

[4] Krifka et al mention these sentences as examples of *characterising sentences*, which they present as a subset of generic sentences. I discuss what characterising sentences are and whether we can use them to define generics in the next section.

While (19b) and (19d) are much like other sentences used as examples of generics nowadays, (19a), (19c) and (19e) are not. Were those that used (19a)-(19e) as examples of generics confused about the scope of the phenomenon they were aiming to describe? Or has work by subsequent authors suggested that only a subset of those sentences mentioned by Krifka et al have a particular set of features that make them worth investigating as a distinctive category? Either way, the assumption that we have a clear and stable grasp of what a generic is and that it merely needs to be invoked by an example or two is unwarranted.

Second, not defining generics means we cannot explain their characteristic features in terms of their defining features, which risks leaving us with an incomplete picture of generics. To see why, consider work on other topics that employs such a strategy. For example, some work in the literature on conditionals attempts to locate the semantic differences between indicative and subjunctive conditionals in the linguistic differences that are used to distinguish them.[5] Similarly, one of the main arguments in my thesis will be that understanding the lack of quantifier expressions in generics, which I argue distinguishes them from other sentences that express generalisations, is an important part of understanding the nature of genericity.

Consequently, I propose that the definitional issue should be tackled head on. I do this by considering two possible ways of defining generics in §2: First, by common *syntactic features*. However, I argue that syntactic features vary too much cross-linguistically, and even within languages, to be of use in defining generics. Instead, I consider whether we can make use of Krifka et al's characterisation of the *varieties* of genericity to define generics. In §3, I argue that this approach isn't promising either because it won't rule out overtly quantified sentences that express generalisations from being generics. In §4, I propose instead that it is the absence of overt quantifier expressions in some sentences that express generalisations which makes them generics. I discuss the case against (§5) and for (§6) thinking that generics are quantificational. I argue that they are, but covertly so. I end by suggesting that genericity in sentences may be a gradable notion, such that some sentences are more generic than others in §7.

2 Defining generics: two options

The generics literature contains much discussion of different characteristics thought to be distinctive of them. For instance, it is often said that generics tolerate exceptions, that they are law-like, associated with dispositions, and resist contextual restriction.[6] However, these features are not usually used to *define* generics. Given that the question about how generics ought to be defined has barely been discussed explicitly, I will consider the main ways in which generics are *introduced* in articles and books about them as possible candidates for definitions.

[5] See Kratzer (1981, 2012).

[6] See Krifka et al (1995) and Mari et al (2012) for comprehensive discussions of different characteristic features of generics.

2.1 Manifestations of generics

Consider again the following sentences I used as examples of generics at the beginning of this chapter:

(1) Nurses are caring

(2) A heatwave causes water shortages

(3) The child learns to read before it learns to write

In English, those sentences described as generics tend to be bare plurals (like (1)), indefinite singulars (like (2)), or definite singulars (like (3)). Bare plurals (BPs) include a noun phrase constituted only by a plural noun without a quantifier or a determiner expression. Indefinite singulars (ISs) include a noun phrase in which the indefinite articles 'a' or 'an' precede a singular noun. Definite Singulars (DSs) include a noun phrase consisting of a singular noun preceded by the definite article 'the'. Given these similarities, we might be tempted to use them to define generics.

However, the common syntactic forms generics take in English cannot be used to define them. For one, not all BPs, ISs and DSs are generics. Take for instance:

(4) Nurses helped those involved in the accident

(5) A heatwave led to an increase in sales of fans

(6) The child learnt to read by looking at packaging

While (4)-(6) are a BP, IS, and DS respectively, they do not receive generic, but existential interpretations. (4)-(6) are not about nurses, heatwaves, and children *in general*, but about *particular* nurses, heatwaves, and children. In each case, the sentences assert that some nurses, heatwaves, and children exist and then predicate some further property of them, namely that they helped those involved in the accident, led to an increase in fans, or learnt to read by looking at packaging. [7]

Even if we could apply some further tests to weed out sentences like (4)-(6), the syntactic definition of generics would be limited to their manifestation in English. [8] In other languages, generics manifest as different kinds of sentences. For example, in languages which do not have bare nouns, like French, generics can manifest as plural definites.

(7) Les baleines sont en voie de disparition [9] [Whales are endangered]

The diversity of generic sentences, cross-linguistically but even just within English, gives us good reason to look for a definition of generics that goes beyond focussing on their syntactic manifestations.

---

[7] See Krifka (2003) on what determines whether a given sentence receives an existential or generic reading, or in fact both.
[8] Krifka et al (1995:8-14) discuss several. See Rosola (2021) for a discussion of these tests, and Carlson (2008) as well as Greenberg (2004) for criticism of one influential test, according to which generics resist contextual restriction.
[9] Example given by Mari et al. (2012:11).

2.2 Varieties of generics

Instead of defining generics in virtue of shared syntactic features, we might want to use Krifka et al's discussion of what they call the 'varieties of genericity' (1995:5). Although, they do not themselves speak of their discussion as providing a definition of generics, their account is routinely used to introduce what generics are and so I take it that this account is the closest to an accepted definition of generics in the literature.[10]

Krifka et al propose that generics are sentences that exhibit genericity and then give an account of how genericity can arise. Either it can arise as a result of the sentence containing *a kind-referring NP,* or it can arise as a result of the sentence being a *characterising sentence*.

2.2.1 Kind-referring NPs

The first way in which genericity can arise is via kind-reference. Specifically, Krifka et al argue that sentences can be generics in virtue of containing a kind-referring NP.

(8) Tea was popularised in China during the Tang period

The idea is that (8) ascribes a property not to individual units of tea, but to tea *as a kind*. Sentences like (8) don't characterise the kind tea by reporting on what is true of individual units of tea. Instead, (8) predicates something of the kind tea directly. In the same way, properties like being extinct or being discovered can be ascribed to kinds like woolly mammoths, though not to individual woolly mammoths.

2.2.2 Characterising sentences

The second variety of genericity is displayed by the following sentences:

(9) Beans cook quickly in a pressure cooker

(10) Boars have bristly hair

(9)-(10) characterise kinds by generalising about their members. Krifka et al say that they characterise kinds by ascribing 'regularities which transcend particular facts' (1995:3). (9)-(10) don't say things about particular beans or particular boars. Instead, they say something about the kinds in question by generalising over particular bean-cooking times and the fur of particular boars.

The resulting picture then is as follows. For Krifka et al, genericity can arise either at the level of the NP, or at the sentence level, through the way properties are ascribed. There are worthwhile questions to be asked about Krifka et al's overall framework. For instance, why should we think that there is anything important that sentences like (8) on the one hand, and (9) and (10) on the other, have in common that makes it such that they should both come under the banner of generics? Both sentences like (8) and

[10] For instance, see Lazaridou-Chatzigoga (2019), Leslie (2008), Nickel (2016), Pelletier (2009).

sentences like (9) may be thought to characterise kinds, tea, and beans respectively, but so do many other sentences which are not categorised as generics. I will return to this point shortly.

In addition, each of the core notions Krifka et al use is left at an intuitive level. While it is clear to me how 'tea' in (8) refers to the kind tea and not individual units of it, the same is not the case for many of the other examples used by Krifka et al. Take for instance: [11]

(11) Potatoes are served whole or mashed as a cooked vegetable

Krifka et al say that (11) is both a characterising sentence and contains a kind-referring NP. However, it is unclear to me why we should think that 'potatoes' in (11) refers to the kind potatoes and not individual, if arbitrary, potatoes. With regards to characterising sentences, the notion that they abstract from particular events, while intuitively plausible, demands an informative account of what this abstraction consists in. I won't pursue these questions here. Instead, I will argue that even if the relevant notions could be elucidated, they cannot provide the basis for a definition of generics.

3 The problem

The problem is this: many sentences that I will argue should be distinguished from generics contain kind-referring NPs or are characterising sentences. The examples of generics that I have used so far have been sentences that differ from generalisations like:

(12) Most nurses are caring

(13) A heatwave always causes water shortages

(14) All Children learn to read before they learn to write

(12)-(14) differ from the sentences I have so far used as examples of generics in that they contain quantifier expressions. Quantifier expressions are expressions that indicate the presence of a quantifier at the level of logical form. For example, the quantifier expression 'all' indicates the presence of a universal quantifier. For readability, I will call sentences like (12)-(14) 'overtly quantified generalisations' and not 'overtly quantified sentences that express generalisations' in this chapter, even though I take the generalisations to be what sentences can express, a distinction that will become more important in later chapters. Overtly quantified generalisations are routinely contrasted with generics.[12] However, overtly quantified generalisations can both be characterising sentences and contain kind-referring NPs. Hence, if generics are to be distinguished from overtly quantified generalisations, they cannot be defined using Krifka et al's account of the varieties of genericity.

---

[11] Example given by Krifka et al (1995:3).

[12] I am not calling these generalisations 'quantified generalisations' because, as I argue in §6 of this chapter, generics are quantificational - the quantifier they contain is just covert.

3.1 Why distinguish generics from overtly quantified generalisations?

A number of features are thought to distinguish generics from other generalisations. However, I'm going to focus on one feature that I think is particularly important. Generics should be distinguished from overtly quantified generalisations because they can be used in different ways.

Famously, generics can be used to convey a large variety of types of generalisations.[13] Generics can be used to make claims about what most members of a kind are disposed to, what they are actually like, how they compare to members of other kinds, or even what they should be like. I investigate this use variability further in Chapters 2 and 3, but for now, note that overtly quantified generalisations are not as flexible.

Quantifiers can be, and often are, restricted, either contextually or otherwise.[14] Consequently, the contents expressed by sentences in which they occur, including generalisations, can vary, but the variability this restriction gives rise to is typically not as radical in overtly quantified generalisation as in generics. For instance, while the 'Most' in (12) can be contextually restricted, say to quantify only over nurses in a particular city, the number and type of claims (12) can be used to make are much more limited than that which (1) can be used to make. While (1) can be used to convey that nurses are more caring than those in a salient comparison class, or even that they should be, (14) cannot be used to make such claims, at least not semantically.[15] [16] So, if we agree that there is good reason to distinguish generics from overtly quantified generalisations, then the following should worry us.

3.2 Overgeneration

Krifka et al's criteria overgenerate generics because overtly quantified generalisations can feature both kind-referring NPs and be characterising sentences.

Each of (12)-(14) contain NPs that can be understood as reference to kinds: nurses, heatwaves, and children. (12)-(14) are not about individual nurses, heatwaves, or children any more than the corresponding generics (1)-(3) are. Hence, insofar as we can make sense of (1)-(3) being about kinds, so are (12)-(14). Consequently, if having a kind-referring NP is sufficient for being a generic then (12)-(14) are generics.

---

[13] I will be talking about 'use variability' instead of 'truth-conditional variability' in this chapter. The reason for this will become apparent in Chapter 2.

[14] For different approaches to how such restriction arises, see Barwise and Cooper (1981), von Fintel (1994), and Stanley and Gendler Szabo (2000).

[15] I discuss whether the different types of generalisations conveyed by generics are conveyed semantically or pragmatically in Chapter 3.

[16] For other differences between generics and overtly quantified generalisations, see Mari et al (2012).

Each of (12)-(14) are also characterising sentences. They all predicate properties not in virtue of any particular individuals or events. Instead, (12)-(14) ascribe the properties in question in virtue of some more general regularities, as characterising sentences do.

Krifka et al say a few things about the relationship between characterising sentences and overtly quantified generalisations.[17] On the one hand, Krifka et al introduce characterising sentences as being 'opposed to explicitly quantified general sentences such as […] *Each potato in this room was grown in Alberta*' (1995:3 fn.2). However, they then also mention several adverbial quantifier expressions which 'enforce a characterising reading', such as 'usually, typically, always' (1995:7). Further, they explicitly say that characterising sentences 'put no limitations on what kinds of NPs may occur in them', including 'quantified NPs' (1995:8). This seems to suggest that some but not all types of overtly quantified generalisations can be characterising sentences.

(15) Each potato in this room was grown in Alberta.

(15), the sentence Krifka et al use to illustrate the difference between characterising sentences and some overtly quantified sentences, is a generalisation and features a quantifier expression, namely 'each'. One might think that (15) is not a characterising sentence because it ascribes the predicate of having been grown in Alberta to particular individuals: those potatoes in the room.

It is an interesting question to ask which overtly quantified generalisations ascribe properties of particular individuals. Are natural language quantifiers intrinsically abstract or simply a useful means by which to abbreviate talk about particulars? [18] Krifka et al's portrayal of characterising sentences assumes that natural language quantifiers, or at least 'each', is not abstract in this sense.

Independently of this wider question, it seems clear to me that while it may not be the case that all overtly quantified generalisations are characterising sentences, some, including adverbial quantifiers like 'usually', 'typically', and 'always', seem to be concerned with particular events or individuals as little as characterising sentences that lack quantifier expressions. Thus, insofar as the category of characterising sentences has any applicability, there is no reason to believe that they exclude all overtly quantified generalisations.

What this means is that if generics are defined as being either characterising sentences or sentences containing kind-referring NPs, then some overtly quantified generalisations will turn out to be generics. In fact, all the examples of overtly quantified generalisations I have used so far, possibly except for (15), are characterising sentences that also contain a kind-referring NP. Consequently, even if we were interested in only one subclass of generics, i.e., only characterising sentences or only sentences containing

[17] Thanks to Martina Rosola for discussion on this issue. See Rosola (2021) for a detailed reconstruction of Krifka et al's characterisation of generics.

[18] For instance, see Wittgenstein (1922), who analyses universal and existential quantifiers as equivalent to conjunctions and disjunctions about particulars.

reference to kinds in the NP, we cannot draw a clear distinction between generics and overtly quantified generalisations with reference to whether the NP is kind-referring or the sentence characterising. Being a characterising sentence or having a kind-referring NP are not individually or jointly sufficient conditions of being a generic. Instead, I propose that the extent to which generics contain quantifier expressions is a more useful basis for distinguishing between generic and non-generic generalisations.

4 Defining generics as unmarked

Instead of trying to define generics in terms of the source of their genericity or the ways in which they manifest, we should make direct use of the differences between them and overtly quantified generalisations in defining them. Generics are generalisations that do not contain quantifier expressions.

4.1 Markedness

Though usually not used to explicitly define generics, the idea that generics are unmarked plays an important role in theorising about them and is often mentioned when introducing them.[19] According to Östen Dahl, 'it turns out that the most general statement that can be made about generics is indeed that they are not overtly marked for tense and aspect, or alternatively, that they employ the least marked tense-aspect choice in the language' (1995:415). Sarah-Jane Leslie also uses unmarkedness as a way of distinguishing generics from other generalisations: '[…] generic sentences contain no such marking – there is no item in them that specifies what sort of generalization is here being expressed' (2012:19).

But what is markedness? Dimitra Lazaridou-Chatzigoga (2019) points to Martin Haspelmath's (2006) definition of *formal markedness* as overt coding by affixes or auxiliary words.[20] The idea is that expressions or sentences contain overt realisations of operators or things like tense or aspect at the surface level. In overtly quantified generalisations the overt codings are *quantifier expressions*.

Claims about unmarkedness are also often used as a premise for various arguments about what generics are like. For example, Leslie (2012, 2008) uses the assumption that generics are unmarked to argue that they express cognitively primitive, default generations, whereas David Liebesman (2011) argues that generics are not quantificational on the basis that they do not contain quantifier expressions.

(16) All okapis are solitary

(17) Okapis are solitary

Whereas (16) contains 'all' to mark the presence of a universal quantifier, (17) contains no such marker. However, the large variety of constructions that can give rise to generic interpretations in different natural languages makes it hard to argue that generics are completely unmarked in this sense.

---

[19] For instance, see Lazaridou-Chatzigoga et al (2017), Lerner & Leslie (2016), Leslie (2014), and Sterken (2015a).
[20] Haspelmath (2006:3)

4.2 Cross-linguistic data

Even in English, to say that generics should be defined in virtue of being completely unmarked would be an oversimplification. While bare plurals don't contain any affixes or auxiliary words to mark the presence of a generic operator, indefinite singulars (IS) and definite singulars (DS) contain 'a' or 'an' and 'the'. To say that IS and DS generics are completely unmarked seems to presume a certain interpretation of them, namely one according to which those expressions are not quantifier expressions nor overt realisations of any other operators.

A number of natural languages contain other notable markers. Dahl and Krifka et al discuss the case of Swahili. Swahili features a verbal prefix 'hu-' whose presence in a sentence indicates that it is a characterising sentence. However, its absence doesn't mean that the sentence in question is not a generic. In fact, the same meaning can be expressed by two different sentences, one containing 'hu-' and one lacking it.[21]

Similarly, Dimitra Lazaridou-Chatzigoga et al (2019) discuss the case of Greek as an example of a language in which generics are as formally marked as the non-generic analogues. In Greek, generics are most often expressed using definite plurals, sentences of the form *The Fs are G*. Definite plurals in Greek are ambiguous between generic and specific readings. What these examples show is that there are natural languages in which some generics are marked and that even in English, markedness or a lack thereof is much more complicated than my naïve definition made it out to be.

However, I still take there to be something valuable about distinguishing generics from overtly quantified generalisations in terms of markedness. For one, the languages in question either only mark a subset of generics and/or involve markers that also feature in non-generic sentences. In no known language is there a unique generic marker that appears in all and only generic sentences. Therefore, the unmarkedness claim may be better understood as the idea that generics don't contain any *dedicated* marker. Leslie puts the idea as follows: 'While languages differ dramatically in how generic claims are syntactically manifested, one does not find phonological forms that are dedicated to the expression of generic claims – generics are always unmarked in this way' (2012:19). Further, in those natural languages which include generic markers, the markers indicate aspect and/or tense. For example, the 'hu-' prefix in Swahili attaches to the verb, indicating that the property is true of the subject in virtue of a habit or other type of regularity, not the presence of a particular quantifier that makes the truth-conditions of the sentence transparent. This is significant because the resulting sentences seem to, at least in English, display more use variability than overtly quantified generalisations.

[21] Dahl (1995) actually discusses a large range natural languages including, but not limited to Swahili. Specifically, he mentions Wolof, Kalaallisut, and Maori as natural languages which contain overt markers for generics. More recently, Taofeeq (2018) has argued that generic sentences are marked in Yoruba.

4.3 Use variability

Whereas overtly quantified generalisations contain quantifier expressions that indicate how many members of a kind need to have the property in question in order to be true, generalisations without quantifier expressions don't have this feature. Sentences like *All Fs are G* wear their truth-conditions on their sleeves. They are true just in case all Fs are G.

In contrast, generic sentences don't have any quantifier expressions that, on the surface level, indicate how many Fs need to have the property in question for the generic to be true. The fact that non-overtly quantified generalisations don't wear their truth-conditions on their sleeves leads to various features being especially pronounced in those sentences, notably use variability.[22]

The source of this variability can be thought of as arising from the fact that the interpretations of generics are not constrained by the presence of a quantifier expression. I will discuss use variability in Chapter 2 and argue that the lack of quantifier expressions in generics lead them to express non-specific generalisations which are more versatile in their use than sentences expressing other generalisations in Chapter 3.

I do, however, think that there are important continuities between generics and overtly quantified generalisations in terms of their use profile: some overtly quantified generalisations can be used in a much more flexible manner than others. For example, sentences featuring quantifier expressions like *typically, generally, many* and *normally* tend to display more use variability than sentences containing *all* or *most*. Thus, I am not opposed to the idea that variability and, with it, genericity (mediated through different types of marking) comes in degrees. I return to this idea in §7 of this chapter.

4.4 Generics as bare generalisations

If the absence of quantifier expressions has this distinctive effect that informs much theorising about generics, then why not use it to define generics, explicitly in contrast to overtly quantified generalisations? On this proposal, generics would be generalisations about kinds that lack overt quantifier expressions. [23]

However, the proposal would rule out some sentences that are usually assumed to be generics. Take:

(8) Tea was popularised in China during the Tang period

(8) is, under the new proposal, not a generic because it is not a generalisation and only generalisations can be generics. I think that this need not be a disadvantage of the proposal. The literature of generics focuses

---

[22] I discuss why the variability in question is in use and not in the truth-conditions of generics in Chapters 2 and 3.

[23] In order not to undergenerate, the proposal cannot be that generics may contain *no* quantifier expressions at all. Consider:
(18) Cats distrust **every** person they encounter
(19) Old people remember **some** childhood friends
Even though (18) and (19) contain quantifier expressions, they nevertheless seem to be generics. This means that the prohibition on quantifier expressions in generics needs to be a prohibition on quantifier expressions only insofar as they occur in particular syntactic or semantic structures. I return to one possible articulation of such a proposal in §3.3.

overwhelmingly on characterising sentences and provides analyses of them that do not extend to sentences like (8).[24] In fact, this may be seen as a symptom of the fact that sentences like (8) should not be counted as generics.

It is true that both sentences like (8) and generics are in some sense about kinds, either by ascribing properties to a kind directly or ascribing them to their members. However, there are also many other types of sentences that can be about kinds, such as overtly quantified generalisations or sentences that contain references to kinds in places other than the NP. Hence, while it is useful to discuss the similarities between different types of sentences that are about kinds, they need not all count as generics.

5 Quantification in generics pt. 1: The case against

The preceding section might make it look as though I am suggesting that generics are not quantificational, whereas other generalisations are, and that this is what distinguishes the two. That's not the case. In the rest of this chapter, I argue that generics are quantificational; it's just that the quantifier they contain is covert.

The generics literature consists of a majority who believe that generics are quantificational, and a minority who believe that they are not. The former believe that generics are much like overtly quantified generalisations with regards to their logical form, whereas the latter liken generics to direct kind predications.[25]

Are generics more like overtly quantified generalisations, such as:

(20) All/some/most/many okapis are solitary

(21) Okapis are typically/usually solitary

Or are they more like simple sentences consisting of a subject and a predicate with no quantifier expression present, like:

(22) Olivia is solitary

I will be arguing in favour of a quantificational approach, both because there are issues with non-quantificational approaches and because I believe that the quantificational approach is independently motivated. I'll start by reviewing the case for non-quantificational approaches.

5.1 Non-quantificational approaches

Liebesman gives two core reasons against the quantificational approach. His first argument is that given that sentences like (8), *Tea was popularised in China during the Tang period*, are kind-predications, this should

[24] Some exceptions are Liebesman (2011) and Teichman (2016).
[25] Carlson (1977), Liebesman (2011) and Teichman (2016)

also be the case for other generics. I have argued that it isn't clear that sentences like (8) should count as generics and so will leave his first argument aside here.[26] His second argument is that the lack of markedness in generics is a core reason against the quantificational approach. He diagnoses a reason for why others have had little to no success in giving a plausible semantic analysis of Gen: 'I think that the intractability of Gen is unsurprising. On my view, Gen has proven intractable for a very simple reason: it doesn't exist' (2011:411).

Liebesman thinks that generics receive generic interpretations, not because they are generalisations about members of kinds, but because they are about kinds themselves. On the kind predication approach, generics don't contain quantifiers that tell us how many members of a kind need to have some property for the generic to be true. Instead, a generic predicates some property of the kind itself. For example, according to the kind predication approach, (1) predicates being caring of the kind nurse, not of individual nurses.

5.2 Category mistake

An initial problem with this approach is that it leaves unclear in virtue of what many generics are true.[27] It doesn't seem like kinds can be caring, lead to increases in the sales of fans, or learn how to read. Those are just not the types of properties that kinds can have. Kind predication theorists must thus give an account of how it is that kinds can have such properties to explain how the relevant generics can be true.

(23) Tigers have stripes

Liebesman provides such an explanation. Instead of accepting that being striped is a property that cannot apply to kinds, Liebesman argues that it can. He says that though it seems counterintuitive, the kind tiger is actually striped in the same way individual tigers are.

Liebesman uses an analogy to explain how this can be the case. He notes that a table can touch a wall in virtue of the fact that some its parts do. According to Liebesman, the table inherits the property of touching a wall from its parts. This is true with regards to tables, but I don't see how this gives us reason to think that that is the case with the kind tiger and the property of being striped. David Liebesman and Ofra Magidor (2017) develop a wider notion of *property inheritance* that is intended to bolster this argument, though Liebesman cautions that 'given the multiplicity of properties and the multiplicity of ways kinds inherit properties from their members, no fully general account of inheritance will be forthcoming' (2011:420).

---

[26] Liebesman thinks that a uniform treatment of generics is motivated especially by sentences that include direct kind attributions alongside characterising attributions. For example, as in *Mosquitoes are both widespread and irritating*. I believe that quantificational approaches can account for such sentences, something I briefly discuss in Chapter 3. In addition, sentences like *Mosquitoes are widespread and always irritating* don't provide good reason for thinking that *Mosquitoes are always irritating* should not be given a quantificational analysis. See Sterken (2016) for discussion.

[27] This concern is discussed by Carlson (1977) and Nickel (2006).

Although Liebesman's inability to give a systematic account of property inheritance mirrors everyone else's struggles to give uniform truth-conditions for generic sentences, something I discuss in the next chapter, I think that his predicament is worse. The issue for everyone else is that is it difficult to explain *how* something most accept is the case is so, namely that (23) is true. In contrast, Liebesman struggles to explain something that also seems intuitively implausible, namely that kinds can be striped.

6 Quantification in generics pt. 2: The case for

6.1 Multiple readings argument

The reason for why quantificational approaches to generics are standard however, is not just that there are issues with kind-predication accounts, but also that the former is independently motivated. I now turn to the most influential argument in favour of quantificational approaches to generics. Carlson (1989) draws attention to the following kinds of sentences, which are ambiguous between different generic readings:

(24) Typhoons arise in this part of the Pacific

(a) Typhoons are such that they (generally) arise in this part of the Pacific

(b) This part of the Pacific is (generally) such that typhoons arise in it

(24a) and (24b) have different truth conditions. If typhoons regularly arose in the area designated by 'this part of the Pacific', but just as or even more often in other areas, then (24a) would be false, while (24b) would be true. The kind-predication view struggles to explain this ambiguity, while the quantificational approach handles it with ease.

Those in favour of the quantificational approach generally believe that generics are quantificational in virtue of containing a quantifier, commonly called 'Gen'. Consequently, generics are thought to have a tripartite structure, consisting of Gen, a matrix, and a restrictor:

Gen x [F(x)][G(x)]

Proponents of quantificational theories can explain the ambiguity of (24) by saying that it arises because of different mappings of materials onto the matrix and restrictor, respectively.

Gen x [typhoon (x)] [arise in this part of the Pacific (x)] (24a) reading

Gen x [this part of the Pacific (x)] [typhoons arise in (x)] (24b) reading

In this way, whether parts of the sentence are mapped onto the restrictor or the matrix of Gen determines the reading (24) receives. These types of scope ambiguities are a familiar phenomenon that also occur in overtly quantified generalisations.

(25) Typhoons often arise in this part of the Pacific

(a) Typhoons are such that they often arise in this part of the Pacific

(b) This part of the Pacific is such that typhoons often arise in it

(25) can be given the same analysis as (24), namely as a case of scope ambiguity. The quantificational theorist can therefore explain the ambiguity in (25) in a straightforward manner that also provides further support to their claim that generics are quantificational in virtue of the similar ambiguities that occur in overtly quantified generalisations.

In contrast, the kind-predication theorist has trouble explaining the ambiguity in (24). Because they do not believe that generics contain quantifiers, and instead, have a simple subject-predicate form, the only generic reading they can account for is this:

(Typhoons) arise in this part of the Pacific

This is the same reading as in (24a), but the kind-predication theorist cannot, at least not straightforwardly, account for the (24b) reading. According to their theory, the subject term denotes the kind and the predicate the property ascribed to the kind. Because the subject term in (24) clearly is 'Typhoons' and the predicate 'arise in this part of the Pacific' there are no grounds on which to explain the availability of both (24a) and (24b) as a type of scope ambiguity.

In response, Liebesman argues that while (24) is ambiguous, it is not ambiguous between two generic readings. Instead, he argues that the (24b) reading is an existential one akin to *Some typhoons arise in this part of the Pacific.* As I explained in §2 of this chapter, bare plurals can give rise to both existential and generic readings, so it is not strange for Liebesman to claim that some bare plural sentences may have an existential interpretation.

If (24) is merely ambiguous between a generic and an existential reading, then there is no special problem for the kind-predication theorist. 'Typhoons' in this case is ambiguous between referring to a kind and referring to particular typhoons. On the generic reading, the sentence is true when it's true of the kind typhoons, that they arise in this part of the Pacific, whereas on the existential reading, the property of arising in this part of the Pacific is predicated of (some) particular typhoons.

Liebesman motivates his view that (24) has an existential reading by pointing out that it is upwards entailing. What this means is that it's possible to substitute 'typhoons' for a more inclusive term, like 'storms' without a change in the truth-value of the sentence.[28]

(26) (a) Birds fly

(b) Animals fly

[28] This is suggested as a test to tell between generic and existential readings in Krifka et al (1995:13-14).

(27) (a) Birds are sitting on the power line

(b) Animals are sitting on the power line

Whereas (27), the sentence receiving an existential interpretation, retains its truth-value after substitution, (26), the sentence receiving a generic interpretation, doesn't.

Now consider:

(24) Typhoons arise in this part of the Pacific

(28) Storms arise in this part of the Pacific

Liebesman says that whenever (24) is true on some reading, (28) is also true and thus, at least one reading of (24) must be existential. I agree with Liebesman, but don't think that this shows what he sets out to show.

While the upward entailment test shows that (24) has an existential reading, it does not show that it only has *one* generic reading. It does seem to me that (24) has several generic readings, including the (24b) reading, which is not the existential reading Liebesman has in mind. While the existential reading would be true if one typhoon had arisen in this part of the Pacific, (24b) would not be true under these circumstances. Therefore, it seems to me that while (24) might have some existential reading, it does also have several generic ones and that as such, the original issue of explaining the ambiguity between generic readings for the kind-predication theorist remains.[29]

The multiple readings argument has been the main reason for why the quantificational approach has been so dominant.[30] However, there are other options on the table. Ariel Cohen (2012) and John Collins (2015) argue that pragmatic and/or psychological processes are responsible for the quantificational force of generics. For example, Cohen argues that generics have the same logical form as kind predictions but get quantificational force as a result of pragmatic processes, specifically a process of 'reinterpretation' involving predicate transfer.

However, given that they face similar issues to the kind-predication approach precisely because they lack the resources that come with a quantificational approach, I will assume a quantificational view for the remainder of this thesis. Some of the advantages of the account I present in Chapter 3 also provide indirect arguments for a quantificational view. For example, I argue that the type of restricted existential quantifier generics contain, interacts with negation in such a way so as to explain the truth-conditions of negated generics.

---

[29] See Carlson (1989) for a more detailed argument to this effect.

[30] For instance, see Sterken (2016) for a similar but distinct argument which also explicitly targets other non-quantificational views like those of Cohen (2012) and Collins (2015).

7 Unmarkedness and covert quantifiers

In §4 of this chapter, I suggested that what makes sentences generics is their *lack* of quantifier expression. However, in §5 and §6, I argued that generics are indeed quantificational. So how can generics both be quantificational and not contain quantifier expressions?

7.1 Convert quantification

The answer is: through covert quantification. The status quo in the literature is that generics contain a covert quantifier, which is not phonologically realised. So, while generic generalisations do contain a quantifier, its presence is not indicated at the surface level by quantifier expressions. Covert quantification is not a novel idea and has precedent in a number of theories which posit logical structures that are not phonologically realised.[31] The idea, then, is that generics contain covert quantificational structures, which can be theorised about in various ways.

For example, many think that Gen is a kind of restricted universal quantifier.[32] Further, most think that Gen has more in common with A-quantifiers referred to by quantifier expressions like *normally, often,* or *usually*, than with D-quantifiers, referred to by quantifier expressions like *the*, *many,* or *most*. [33] [34]

However, given some of my remarks on markedness in §4, the matter is not so simple. Some sentences that are counted as generics do contain expressions like 'a' and 'the' that are widely assumed to be quantifier expressions, at least when they feature in non-generic sentences. This is the case especially in languages other than English, which do not allow bare noun phrases. How can we understand generics as being defined in virtue of their lack of markedness if many generics, in English and in other languages, are very much marked?

My response is twofold: first, it is important to distinguish between different types of markedness, as not all markedness involves quantifier expressions; second, it may be plausible to think of genericity as a feature that comes in degrees corresponding to the extent to which generalisations are marked.

With regards to the first point, I noted in §4 that it is important to distinguish between sentences that are marked in *some way* from those that are marked using quantifier expressions. At least considering generics in sentences other than English, the way they are marked is often not through straightforward quantifier

---

[31] Some other examples include:
-covert modals (see Bhatt 2008 for an overview)
-distributivity operators in the semantics of plurality, see Link (1987), Roberts (1989) and Lasersohn (1998))
-exhaustivity operators in work on scalar implicatures, see Spector (2016)
-existential closure that is thought to occur in a variety of ways: in descriptions (Graff Fara 2015), over events Davidson (1967), or posited as part of dynamic semantics (Heim 1982). I argue that generics get their quantificational force as a result of existential closure in Chapter 3.

[32] I discuss the *normalcy account* family, the historically dominant approach on generics, in Chapters 2 and 3.

[33] D-quantifiers are marked by determiner expressions. In contrast, A-quantifiers are marked by expressions that are adverbs, auxiliaries, affixes, or argument-structure adjusters.

[34] See Cohen 2001b for an overview of the reasons for why this is the case.

expressions, but through e.g., markers of tense. In English too, while definite and indefinite singulars do contain 'a'/'an' or 'the', these are not expressions that, in generic sentences, quantify over concrete individuals as other quantifier expressions can. So, it is important to clarify that generics need not be wholly unmarked, but that the way they are marked tends not to be through a straightforward appearance of quantifier expressions that make transparent the conditions under which they are true or false.

However, it may also be worth acknowledging that markedness, including through quantifier expressions, need not be an all or nothing matter. Instead, one way of accounting for differing degrees of markedness is by thinking of genericity as something that arises in degrees. The more marked a generalisation the less generic it is and conversely, the less marked a generalisation the more generic it is.

To test this hypothesis, we would need to decide which features of genericity are thought to be particularly characteristic and find out whether these are more apparent in less marked generalisations. One other feature, which I discussed in Chapter 4, is use variability.[35] I won't be able to conduct a comprehensive consideration of this matter here, particularly a cross-linguistic one. However, just looking at English, there are initial reasons for thinking that IS and ID generics display less use variability than BP generics.

7.2 Markedness in IS and DS generics

There is a literature that suggests that non-bare plurals receive more restricted readings than BP generics.

7.2.1 IS generics

(29) (a) Dogs bark

(b) A dog barks

(c) The dog barks

At first sight, (29a), (29b), and (29c) seem relatively similar to one another. They all express generic generalisations about dogs, namely that they bark. However, they also differ from one another. John Lawler (1973) argues that IS generics are in some sense definitional, whereas bare plural ones are inductive. This contrast is illustrated by the following:[36]

(30) (a) Madrigals are polyphonic

(b) A madrigal is polyphonic

[35] In Chapter 4, I will discuss work by Lazaridou-Chatzigoga et al (2019) which compared inference patterns between generic and overtly quantified generalisations and found (i) no clear dividing line between generics and overtly quantified generalisations and (ii) important differences between different types of generic sentences. Such work may further strengthen the idea that genericity, insofar as it is defined in opposition to quantified generalisations, is a matter of degree.

[36] Examples are by Lawler (1973).

(31) (a) Madrigals are popular

(b) *A madrigal is popular

The intuition reported by Lawler is that while (30a) and (30b) are both felicitous, (31a) is felicitous, but (32b) is not. According to Lawler (1973), this is because IS generics are only felicitous when they predicate properties that are intuitively essential to the kind. Being polyphonic is essential to being a madrigal while being popular isn't, hence (30b) is felicitous and (31b) isn't.

Further, Noel Burton Roberts (1976) argues that IS generics have a stronger normative force than BP generics and Cohen (2001) interprets this force as IS generics making claims about rules and regulations following Carlson (1995), not regularities, which he takes BP generics to be in the business of making. [37]

7.2.2 DS generics

It's not just indefinite singulars that work differently from bare plurals. Definite singulars have also been thought to have more restricted interpretations than bare plurals. Consider the following examples[38]:

(33) (a) The coke bottle has a narrow neck

(b) *The green bottle has a narrow neck

(c) Coke bottles/green bottles have a narrow neck

While (33a) seems felicitous, (33b), on a generic reading, does not. This may be because while coke bottles are what Krifka et al call 'an established kind' (1995:11), green bottles are less so. Note also that the generic in a bare plural form is fine regardless of whether the kind is established or not.

So, at least at first sight and in English, it seems as if IS and ID generics are both more marked than BP generics and display less use variability than them, hence providing some initial evidence in favour of my hypothesis. Consequently, I will focus on BP generics as paradigmatic generics for the remainder of this thesis.

8 Conclusion

Generics cannot be defined in terms of being characterising sentences or containing kind-referring noun phrases. This is because sentences that we want to distinguish generics from, overtly generalisations, can contain reference to kinds and be characterising sentences. Hence, if we want to distinguish generics from overtly quantified generalisations, we need to find a different criterion.

---

[37] Yael Greenberg (2004) discusses these differences in detail and accommodates them in her account of generics by positing that BP and IS generics quantify over different domains. According to Greenberg, both BP and IS generics contain Gen, a vague, restricted universal quantifier. However, the domains over which Gen ranges depends on the syntactic form of the sentence resulting in the availability of different readings for BP and IS generics. Also see Nickel (2016) for a detailed discussion of Greenberg's approach.

[38] Example discussed in Krifka et al 1995. Contrast is given by Carlson (1977), who attributes it to Barbara Partee.

I proposed using the lack of surface quantifier expressions in generics to distinguish the two. Doing so allows us to focus on a distinguishing feature of genericity: use variability that generics exhibit as a result of not containing explicit quantifier expressions. This criterion raises interesting but thorny questions about the extent to which generics are unmarked. I suggested that this is a virtue, not a vice as it opens the door to thinking of genericity in degrees. Focusing on forms and degrees of markedness allows us to distinguish generics and overtly quantified generalisations from one another, while also accounting for the continuities between them.

Chapter 2: Non-specific generics

1 Introduction

Finding and articulating the conditions under which generics are true or false has been the principal undertaking in the literature on generics. 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 the cognitive, political, and moral significance of generics.[39] However, determining the truth-conditions of generics has been difficult because the facts in virtue of which they can be true vary radically. In this chapter, 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 begin by surveying some popular attempts to accommodate truthmaker variability in generics (§2). In §3, I articulate an assumption made implicitly in those approaches, namely that generics express generalisations of a specified quantificational force and flavour and present a counterexample to it. I describe recent discussions of truth-conditional variability across generic tokens (§4) and consider how they may be brought to bear on non-specific uses of generics (§5). I discuss two possible ways of explaining non-specificity, namely that generics uttered in indeterminate contexts express no propositions (§6) or that they express several (§7). I argue that understanding non-specificity as a by-product of meta-semantic phenomena like context-sensitivity or semantic incompleteness invites a host of issues and limits the explanatory role of non-specificity.(§8) Instead, I propose that generics, by default, express non-specific generalisations, a view I then go on to develop in Chapter 3.

1.1 Non-specificity

Non-specificity is going to play a big role in this chapter and the ones to follow, so let me start by briefly clarifying what I mean in ascribing it to generics. All sentences that express 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.[40] However, in this chapter, I will argue that generics are non-specific in a further sense: they need not specify the quantificational force or flavour of the link between members of a kind and the property in question. For example, while *Most dogs bark* specifies a

[39] For instance, see:
- Leslie (2007, 2008, 2012) on generics as articulating cognitively basic default generalisations
- Leslie (2015, 2017), McConnell-Ginet (2012), Haslanger (2011, 2014), Saul (2017), Ritchie (2019), and Lemeire (2020) on generics as vehicles of stereotyping
- Nickel (2010) on natural laws as generic
- Lerner and Leslie (2013) on the role of generics in moral theorising
- Wasserman (2011) on the relationship between generics and dispositions.

[40] Thanks to an anonymous reviewer for Synthese for pointing out this further sense in which generics are non-specific.

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.

2 Truthmaker variability in the wider literature

Although the literature on generics is young, there are plenty of theories about the truth conditions of generics. However, these theories sooner or later struggle with one feature of generics. Consider the following:[41]

(1) Rocking chairs are made from wood

(2) Elks have antlers

(3) Ticks carry Lyme disease

(4) Elms are trees

(5) Girls don't play in puddles

(6) Brazilians are good football players[42]

(1) seems to be true because most rocking chairs are made from wood. In contrast, (2) is true even though most elks don't have antlers. Only a minority of elks, namely bulls over a certain age, grow antlers. (3) is an even more extreme example of this. (3) seems to be true, even though only a small percentage of ticks actually carry Lyme disease. So, while the truth of some generic sentences appears to have something to do with the high incidence of a property among members of a kind, others can be true even though the property is only had by a small minority. It is hard to say in virtue of which kinds of facts (2)-(6) are true for reasons that I will be discussing in this and the next chapter.

However, here are some initial ideas: (2) might be true in virtue of the fact that antlers are typical physical characteristics of elks. (3) might be true because ticks are capable of and disposed to carrying Lyme disease. (4) is true in virtue of what it is to be an elm, namely a type of tree. If (5) was true, it would be in virtue of expressing a normative rule according to which girls ought not to play in puddles. Lastly, (6) might be true because good footballers are disproportionately often Brazilian, even though most Brazilians are not good footballers. The types of facts that can make true (2)-(6) are not facts about what is the case for most elks, ticks, girls, or Brazilians.

This variability in what can make true different generic sentences is commonly conceptualised as variability in the *truth-conditions of generics*. On this way of thinking, the types of facts that can make true (2)-(6) vary just because the conditions under which (2)-(6) are true or false themselves vary. I will argue in the next chapter, that the truthmaker variability is not in fact a case of truth-conditional variability. So as

[41] (1)-(4) taken from Prasada & Dillingham (2006), (5)-(6) are my own.
[42] By 'football', I mean what North Americans call 'soccer'.

not to prejudice discussion, I will be talking about truthmaker variability for now, which I take to be consistent with all explanations that acknowledge any sort of variability.[43] Truthmaker variability has long been thought of as the main reason for why it has proven so difficult to provide uniform truth-conditions for generics. It's hard to understand generics as expressing any one particular type of generalisations if the types of generalisations they express appear to change. Nevertheless, many have tried to provide semantic theories of generics which can accommodate truthmaker variability.

According to one popular approach, generics are *universal generalisations in disguise*, quantifying over normal members of a kind.[44] According to one rough articulation of this approach, (1) is true iff all normal rocking chairs are wooden. One issue accounts of this kind face is that some generics seem capable of being true even though they ascribe properties that are not had by all normal members of the kind, an instance of the truthmaker variability I described above. For example, (6) seems to be true even though not all normal Brazilians are good at football.[45] As a result, those pursuing the normalcy approach tend to incorporate context-sensitivity as well as relying on domain restriction and a thin notion of normality to be able to account for truthmaker variability.[46]

In contrast, Ariel Cohen (1999a; 1999b) argues that generics are *probability judgements*. 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.[47] 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. By allowing two different kinds of reading, Cohen's account can account for at least some degree of truthmaker variability.

Sarah-Jane Leslie (2007, 2008, 2012) takes a different approach. She distinguishes between different *types* of generics. Instead of trying to find a common type of truthmaker for all generics, she builds her theory around the idea that different generics are made true by different types of facts. Leslie argues that some generics are true in virtue of what's characteristic of their kind, others by what members of a kind are

---

[43] One exception might be Nguyen's (2019) account, according to which generic sentences themselves express propositional radicals and are hence not truth-apt. However, even Nguyen thinks that there is truthmaker variability, it's just that, according to him, the truthmaker doesn't make true generic tokens themselves, but implicitures the generic tokens convey. I discuss Nguyen's account in §5.2 of this chapter.

[44] Theories of this type, albeit with various refinements, have been advocated by for example, Dahl (1975), Delgrande (1987), Asher and Morreau (1995), Eckardt (1999), Krifka et al (1995), Greenberg (2004, 2007), Asher and Pelletier (2012).

[45] See Leslie (2008) for discussion as well as Asher and Pelletier (2012) for a response. I will discuss some of the issues at hand in greater detail in §4 of Chapter 3.

[46] I discuss contextualism in §5.1 of this chapter and Nickel's (2016) normalcy approach in §4 of Chapter 3.

[47] Cohen's notion of a relevant member of a kind is complicated but very roughly, a relevant member of the kind is a member who either has the property or some other salient property. For instance, a generic like *Birds fly* is to be evaluated with respect to the set of birds that either fly or have some other means of locomotion, such as walking, swimming, or taking the bus. This alternatives provision helps account for generics like *Birds lay eggs*, which seem true even though only a minority of birds lay eggs. The set of relevant individuals for this generic only includes birds that lay eggs or have some other means of reproduction, such as mitosis or giving birth to live young.

disposed to be like, and yet others by what most of them are like. For example, (1) is true because most rocking chairs are wooden, (2) is true because, although most elks don't have antlers, it's characteristic of their kind, and (3) is true because ticks, even non-infected ones, have the dangerous (to us) property of being disposed to carry Lyme disease.[48]

3 The specificity assumption and a counterexample

The views discussed so far take different strategies with regards to how they accommodate truthmaker variability across generic types. However, there are also commonalities between them. One assumption these views generally make is that generics express fine-grained generalisations: universal 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.[49]

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 'more likely' 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 disposed to be like, what they 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.

3.1 Seagulls

My friend Aoife and I are discussing things she doesn't like about Aberdeen. So far in our conversation, Aoife has mentioned 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 people's food.' The last part of Aoife's utterance is a generic, but what type of generalisation does it express?

(7) Seagulls swoop down to steal people's food

(a) Most seagulls swoop down to steal people's food

(b) All normal seagulls swoop down to steal people's food

[48] Leslie's approach to generics is much more complex than this in several ways. For Leslie, generics serve a particular role in articulating cognitively primitive generalisations. As a result of this role, the truth-makers she describes are not to be understood as the semantic truth-conditions for generic sentences. Further, generic sentences are sensitive to the nature of counter-instances to the generalisations in ways that track their psychological salience. I discuss Leslie's account of genericity in mental states in Chapter 4 but see Sterken (2015b) and Nickel (2016) for a more general discussion of Leslie's approach, including problems with her proposed truthmakers.

[49] 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 §5.1 of this chapter.

(c) All seagulls are disposed to swoop down to steal people's food

(d) More seagulls than other types of birds swoop down to steal people's food

(e) Swooping down to steal people's food is characteristic of seagulls

I will call (7a)-(7e) '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 (7a)-(7e), are what could make true Aoife's utterance of (7).

First, decide on whichever parsing of (7) strikes you as the most plausible. Now, imagine that that parsing was false, but that another of the candidate generalisations 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 (7) 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 (7a)-(7e)? This doesn't seem plausible to me. We often use generics to make general remarks, 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 (7a)-(7e).[50]

This example doesn't seem far-fetched. 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 rough generalisations that need not be specific in this way. This chapter is about understanding the non-specificity at hand.

One initial thought might be that the context in which Aoife utters (7) matters. Recently, accounts according to which sentences like (7) can convey *different* generalisations depending on the context in which they are uttered have gathered momentum.[51] 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 (7a)-(7e) be what Aoife's utterance expressed.

And indeed, both Sterken (2015a) and Nguyen (2019), who have emphasised flexibility in the truth-conditions of generics across contexts, have discussed what they call 'underspecified' or 'indeterminate'

---

[50] In §3 of Chapter 5, I argue that generic beliefs are evidentially undemanding in a way beliefs in overtly quantified generalisations aren't, which means we can form justified generic beliefs when the evidence we have wouldn't suffice to justify beliefs in more specific generalisations. This can give us an explanation for why Aoife may not have believed any one candidate generalisation, but still believed (7).

[51] See Nickel (2008, 2016), Sterken (2015a), and Nguyen (2019).

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.

4 Token variability

Bernhard Nickel (2008, 2010), Rachel Sterken (2015a), and Anthony Nguyen (2019) have drawn attention to what they think of as truth-conditional variability across generic *tokens*. Sterken and Nguyen distinguish between two different ways in which the generalisations conveyed by generic tokens appear to vary: in their quantificational force or their flavour.[52]

4.1 Quantificational force variability

Nguyen uses the following example:[53]

(8) Lottery tickets are losers

John needs money and asks Sally whether she thinks that he should buy lottery tickets. In response, Sally says: 'No, lottery tickets are losers'. Nguyen argues that in the 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. Donald believes that the lottery is a big scam as all lottery tickets are losers. He signs a contract, stating that he'll give any lottery winner a billion US dollars. His friend, who thinks this was a silly 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 that 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 (8). Take the following example:

(9) Pizzas are in the freezer section

(9) 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 (9) 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 (9) to make a much weaker claim, such that *most* or even just *many* pizzas are in the freezer section.

---

[52] Nguyen and Sterken recognise both types of token variability that I discuss but call it different things. Nguyen speaks of 'statistical' and 'use variety' and Sterken of variability in the 'quantificational force' and 'lexical restrictor' of Gen, respectively.
[53] Example originally given in Sterken (2015:18) and developed by Nguyen (2019:1305).

4.2 Flavour variability

The second type of variability concerns not the quantificational force of generic utterances, but their flavour. Sterken gives several 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.[54] I call this variability 'flavour variability'.

(10) Horses wear horseshoes[55]

Sterken argues that (10) 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. Wild horses tend to live in environments and move in ways which harden their hooves and make additional protection unnecessary. However, domesticated horses, which make up most of the horse population, tend to be kept in environments which soften their hooves and are used for purposes that can lead to injuries. As a result, it is plausible that most horses wear horseshoes.

Tokens of (10) asserted in a discussion of modern horse care are true. However, tokens of (10) asserted contexts in which questions of evolutionary biology are salient, say in a TV programme about evolutionary adaptations, are, according to Sterken, false.[56] Sterken argues that that's because the truth conditions of tokens of (10) vary across contexts. Tokens of (10) asserted in a modern horse care context, say in a conversation between horse breeders, are true iff most horses wear horseshoes. Tokens of (10) asserted in an evolutionary biology context are true iff horses naturally wear horseshoes. Since the former is the case and the latter isn't, tokens of (10) are, Sterken argues, true in some contexts and false in others.[57]

I do think that the examples given by Sterken and Nguyen definitely pick up on a variability in use underwritten by truthmaker variability, although whether that variability is truth-conditional or of a different kind is a matter I return to in the next chapter. At least as the examples are described, I think that the truth-conditions of the generic tokens may be less specific. For example, it seems implausible to me that the shop assistant in the second supermarket must intend to *either* convey the generalisation that most pizzas are in the freezer section, or that most are. Both Sterken and Nguyen discuss non-specific uses of generics even though they 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. However, their views don't commit them to saying how common non-specific uses of generics are and as such, they would be able to

---

[54] Sterken (2015:6-7)

[55] Sterken actually uses the much more widely discussed example of *Dobermanns have floppy ears* introduced by Nickel (2008). It is plausible that it is no longer the case that most dobermans have had their ears clipped. Because the example as presented relies on some prevalence claim being true, I use (10) instead. The considerations are analogous.

[56] Nguyen (2019:1306) gives some further examples of this kind.

[57] 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.

accept widespread non-specificity in generic tokens. 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.

5 Explaining token variability

Several explanations for token variability have been put forward. These responses typically diagnose the difference in meaning between generic tokens as arising from the contexts in which they are uttered. I discuss two such explanations: theories that posit context-sensitivity and theories that posit semantic incompleteness in generics.

5.1 Contextualism

One straightforward way of accounting for the variability data is to argue that generics are context-sensitive. Such an approach has been implemented in various ways.

Sterken uses the observations about token variability to motivate an account, according to which generics are context-sensitive. She argues that the unpronounced generic quantifier, Gen, is context-sensitive. 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 (8), *Lottery tickets are losers,* is true iff all lottery tickets are losers, whereas another token of the same type is true iff most lottery tickets are losers.

Context-sensitivity can be integrated into other extant accounts. For example, normalcy approaches typically include contextual parameters. Contexts of utterance can supply both a modal base and an ordering source.[58] The modal base is the domain of quantification, namely the set of possible worlds that are accessible from the actual world. In some contexts, the worlds are worlds that play out what is statistically likely to happen, others play out what should, yet others what could happen. The ordering source ranks these worlds from most to least normal. For example, if the modal base is statistically likely worlds, the ordering source provides a ranking of these worlds from most to least normal.

[58] Following Kratzer (1977)

By varying features of the context, we get different modal bases and differing ordering sources. In some contexts in which a generic type is tokened, only the statistically likely possible worlds are relevant, so the token has a statistical flavour. In others, only worlds that play out what should be the case are relevant and so the token has a normative flavour. Even in contexts with the same modal base, the ordering source can vary such that what's the most normal statistical world at which to evaluate a token of a type need not be the most normal statistical world at which to evaluate another token of the same type. In this way, we can explain how and why the truth-conditions of generics appear to vary across contexts of utterance.

So, for instance, in some contexts of utterance, a token of (10), *Horses wear horseshoes,* the modal base includes only statistically likely worlds. This token has an absolute prevalence flavour, it says that wearing horseshoes is prevalent in horses. In another context, the modal base comprises teleologically accessible worlds in which everything realises its essence and nothing else happens. It is not the case that in the most normal of those worlds all horses that inhabit it wear horseshoes, hence the token is false.

According to contextualism then, generics are context-sensitive. The truth conditions of generics vary across tokens, depending on certain features of the contexts of utterance. An approach I consider next explains variability not in terms of the generics being context-sensitive, but in terms of them being semantically incomplete.

5.2 Semantic incompleteness

Nguyen (2019) proposes that generics are semantically incomplete. He utilises Kent Bach's notion of propositional radicals (1994, 2001a, 2001b, 2006) 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.[59] According to Bach, the completed content is not the semantic value of the completed utterance but is communicated pragmatically as an impliciture.[60]

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.

---

[59] Bach (2006:436)

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

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 implicitures. For example, in the context of modern horse care, the speaker uttering (10) completes the propositional in such a way so as to convey that most horses wear horseshoes. In contrast, in the TV programme about evolutionary biology, the speaker completes the propositional radical in such a way so as to convey the claim that it is characteristic of horses that they wear horseshoes.

5.3 Token variability and non-specificity

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, whether semantic or pragmatic. 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*, namely one that lacks specific intentions or other ingredients that determine *one specific* generalisation as what is communicated by a generic token.[61] Following Kai von Fintel and Anthony Gillies (2011), who discuss related issues surrounding epistemic modals, I will call such contexts 'indeterminate contexts'.

For Sterken and Nguyen, non-specificity is a by-product of whatever phenomenon they take to be responsible for token-variability. Both place emphasis on the types of *contexts* in which non-specific utterances of generics are made. In describing the type of case, 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) (2015:21-22).

Nguyen also discusses non-specific uses of generics.[62] 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.

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:

---

[61] 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.
[62] Nguyen (2019:1315)

(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 'underspecified' generalisations

I believe that option (a) is implausible. First, picking any one of the candidate generalisations over the others would seem to be an entirely arbitrary matter. On what basis could we pick one candidate generalisation over the others as what the generic token expresses, if not with reference to the context of utterance? 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 chapter.[63] 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 thesis.

How can Sterken and Nguyen account for non-specific generic tokens? Although both acknowledge non-specific uses of generics, neither discusses in much detail what the propositional content of generic tokens uttered in indeterminate contexts is. I will consider two options: that they convey no propositions or that they convey several.

6 Propositional nihilism

If generics rely on some features of the context in which they are uttered for their meaning - for instance through some sort of 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 I am concerned with here, just aren't the kind of sentences that (semantically or through implicitures) 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 will focus on what kind of explanation the semantic incompleteness proponent could develop, but similar considerations apply for contextualist versions. According to Nguyen, the propositional radicals that generics express need to be completed by speaker intentions in order to convey full propositions via implicitures. Thus, if the speakers lack the fine-grained intentions that would complete the propositional radicals, then no full proposition at all is conveyed.

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 and express truth-apt generalisations. If generic tokens themselves only

[63] Although King (2013) explicitly allows for non-specific intentions, I am also unsure of how his account could account for the content to be indeterminate, rather than non-specific. If the speakers intend to convey a range of candidate generalisations, and the audience can coordinate on this, then there is no reason to believe that only one of the candidate expressions is expressed. Alternatively, if the speaker intends to convey a range, but the audience coordinates on only one candidate, then, on the coordination account, there should be no indeterminacy.

express propositional radicals, any truth-apt generalisations must be conveyed via implicitures. However, these implicitures require that speakers have the fine-grained intentions required to complete the propositional radicals in one of several ways. When they don't, *no* propositions are expressed, either semantically or via implicitures.

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. 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.[64]

Second, the tokens in questions don't just seem complete, they also seem to be capable of communicating something true or false. Utterances of *Dogs bark* just say something true, including ones in which their speakers fail to have highly fine-grained intentions. Note that the issue here is not just that Nguyen's proposal entails that generic don't have any truth-conditions, as to those for whom this is a concern semantic incompleteness would never be an attractive proposal in the first place. Instead, the problem is that if the speakers in the relevant contexts lack fine-grained truth-conditions, then no implicitures are generated that could be true or false.

Given the role generic tokens expressing rough generalisations play in our everyday speech as well as generic beliefs, it would seem strange if nothing they said (either semantically or via implicitures) was truth-evaluable and we would need a reason for why they are used in this way. Aoife's utterance is clearly meaningful, but the propositional nihilist has no explanation of how this meaning is communicated either semantically or pragmatically. Consequently, I will explore an alternative response, namely that generic tokens uttered in indeterminate contexts do express propositions, in fact, that they express several.

---

[64] In response, an anonymous reviewer for *Synthese* 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 implicitures, then we need an account of how these are generated by indeterminate intentions. If they are not communicated through implicitures, 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, I think that this response only postpones the issue.

7 Propositional pluralism

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 a result of the sentence's containing context-sensitive expressions (flexible pluralism) or gaps, as in Bach's framework (strong pluralism). This means that pluralism is compatible with both contextualism and semantic incompleteness theories. I will be focusing on flexible pluralism for brevity, 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 proposed such an account for generics. According to flexible pluralism about generics, generic tokens uttered in indeterminate contexts express several generalisations. Generic tokens always express a set of propositions, but this set can contain more than one element when the contexts in which they are uttered are indeterminate. Pluralism does a good job of explaining the 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 of them. 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.

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 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. For instance, Viebahn (2019) argues that propositional pluralism is plausible for a range of expressions, such as quantifiers, gradable adjectives, and modals.

Propositional pluralism has its advantages, but it also comes with some costs that speak in favour of finding an alternative approach. Some of these are general. 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? If not, what does one report by ascribing generic beliefs to others? What norms govern them? Are they only properly asserted if we meet epistemic requirements for all the propositions they convey? 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.

Indeed, 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 their non-specificity. This is a matter I explore in detail in chapters 4-6, but briefly, I believe that propositional pluralism may make generic contents too hard to believe and assert than would be suggested by the actual role they play in cognition and communication.

However, there is also another concern that I have already anticipated in my discussion of Sterken's analysis. I suggested that if one lacks the requisite intentions to convey a fine-grained generalisation, there is no reason to think that any such generalisations are communicated. This thought also applies if intentions aren't the decisive ingredients of contexts. Whatever those ingredients are, if they are lacking in the context of utterance then there is no reason to think that the contents that would be communicated in the presence of the ingredients will nevertheless be communicated.[65]

Applied to the case of generics, if I utter 'Dogs bark' without having any intentions to convey any one non-generic generalisation or without being in a conversation that is about a fine-grained topic, then it just seems unmotivated to think that my utterance communicates any of those candidate generalisations at all because the thing that would partly cause my utterance to have those meanings, fine-grained intentions, are not present. Propositional pluralists will have to make a positive argument for why we should think that speakers can express not just one, but several propositions if they do not possess the fine-grained intentions normally required to successfully use context-sensitive language.

8 Looking for an alternative view

These considerations illustrate that while propositional pluralism has the resources to account for non-specific uses of generics, it also brings with it many complications.

In particular, I suggested that if speakers lack fine-grained communicative intentions, there is no reason to think that their utterances express any fine-grained generalisations, much less several. However, the alternative, namely that *no* truth-apt generalisations at all are conveyed leaves us with questions about why

[65] In response, an anonymous reviewer for Synthese 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 with his strategy would be that it would entail that one would often say more, semantically or via implicitures, by speaking generally than if one was more specific.

we would use generics in the first place. I also noted that thinking of non-specific generics as a phenomenon that arises as a by-product of token variability restricts what such theories can contribute to an understanding of the mental states that are expressed by non-specific generic tokens.

In addition to these complications, there is also a question to be asked about whether there are any good reasons for thinking of non-specific uses of generics as a by-product of some other phenomenon, be it context-sensitivity or semantic completeness, instead of as simply what generics mean. Why should we think that generics, uttered in any contexts, express the same kind of generalisations that could also be expressed using overtly quantified generalisations?

In the next chapter, I therefore want to propose an alternative that avoids the specificity assumption: generics express *one* proposition at a time, just weak, non-specific ones. The basic idea is this: rather than assume that generics express fine-grained generalisations and then explain non-specific uses of them in terms of some other phenomenon, they express less-specific generalisations by default. Generics generalise about members of a kind and attribute a link to them and some property, but they need not say much about the nature of that link.

This account can provide a straightforward explanation of how generics can be used to successfully communicate generalisations in the absence of specific communicative intentions. It can also easily be extended to an account of the kinds of mental states that can be expressed using generics, namely as propositional attitudes towards non-specific generalisations. I will present an account of generics as expressing non-specific generalisations in the next chapter, before extending it to an account of generic beliefs in Chapter 4.

9 Conclusion

In this chapter, I have drawn attention to what I take to be an underappreciated aspect of generics: their non-specificity. Generics are often used in a way that doesn't specify the quantificational force or flavour of the generalisation made. 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. Instead, I proposed that a novel account, according to which generics semantically express non-specific generalisations should be considered. I develop just such an account in the next chapter.

Chapter 3: The Generality Account

1 Introduction

At the end of Chapter 2, I suggested that we should consider generics as semantically expressing non-specific generalisations. In this chapter, I develop such an account. The basic idea behind the account is that generics semantically express non-specific generalisations. Generalisations are always non-specific in one sense; they are not about any particular members of the group generalised about. However, in the previous chapter, I argued that they are non-specific in a further sense, namely in that they need not specify the quantificational force or flavour of the generalisation they express. According to the account I propose, the 'generality account', generics are non-specific in this further sense because they quantify existentially over non-generic generalisations. A generic like *Bats have good hearing* existentially quantifies over non-generic generalisations concerning bats and the property of having good hearing, such as ones ascribing absolute or comparative prevalence, dispositions, norms, or essences. Consequently, it is true as long as at least one of these non-generic generalisations are true. On this account, generics aren't understood as expressing one type of other generalisation, but rather as making *general* generalisations.

I begin the chapter by introducing an invariant version of the generality account according to which generics always quantify over the same set of non-generic generalisations about the kind and property (§2). According to this invariant version of the generality account, all generics are non-specific. In §3, I argue that this existential quantification comes about as a result of a process of existential closure triggered by an absence of quantifier expressions in the generic sentence. In §4, I distinguish this view from its closest existing relative in the generics literature: Bernhard Nickel's (2016) normalcy account. I criticise Nickel's view on the basis that his notion of normalcy is either too restrictive or uninformative. In §5, I describe some of the generality account's advantages, but then respond to a major worry, namely that it overgenerates true generics (§6). In §7, I argue that the generality account can in fact explain why, though generics have weak semantic meaning, they nevertheless give rise to a rich web of pragmatic meaning. This can explain variability in their use while preserving stable truth-conditions for generics, which will form a crucial explanatory posit in the second half of this thesis.

2 An alternative proposal

2.1 The basic idea

My proposal is this: rather than assume that generics 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 Chapter 2, 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.

Though the idea that generics express non-specific, or general, generalisations hasn't been widely discussed in the generics literature, the idea behind it actually seems somewhat obvious.[66] First, the terms 'generics' and 'generality' sound very similar and, it turns out, they are both formed from the Latin 'genus', which means kind.[67] Second, the lack of quantifier expressions in generics has been widely noted, whether to define them, as I have done in Chapter 1, or merely as a characteristic feature. If quantifier expressions in sentences that express generalisations restrict which such generalisations they express, then it doesn't seem far-fetched to think that sentences that lack quantifier expressions express less restricted generalisations.

2.1.1 Specificity in generalisations

Specific generalisations are generalisations that specify the nature of the link between members of a kind and the property in question. For example, these are generalisations about what is normal, what is characteristic, what is prevalent, what dispositions members of the kind have, or even what they ought to be like.[68] According to the generality account, generics encode non-specificity by existentially quantifying over non-generic, more specific generalisations concerning the kind and property in question.[69]

(1) Horses wear horseshoes

For example, (1) existentially quantifies over non-generic generalisations concerning horses and the wearing of horseshoes: that most horses wear horseshoes, that more horses than cows wear horseshoes, that it is a norm that they wear horseshoes, that it is characteristic of horses that they wear horseshoes, and so on. Consequently, (1) is true as long as at least one of the non-generic generalisation about horses wearing horseshoes is true. This is the non-specificity view in its most basic guise.

If generics are non-specific, then it's not surprising that it has been hard to unearth fine-grained truth conditions for them. Much as *My neighbour is in the courtyard* doesn't say anything about which particular neighbour is in the courtyard, generics don't, at least not automatically, say anything about which non-generic generalisations about the kind and property in question is true, just that at least one of them is.

---

[66] The closest theories I am aware of are Nickel (2016), Carlson (2008) who argues that generics quantify existentially over inductive patterns, Greenberg (2004) who argues that they are vague, and Lemeire (2020) who argues that they express disjunctive generalisations. I discuss Nickel's account in some detail in §4 of this chapter.

[67] For further discussion of this see Nickel (2016:14).

[68] I discuss which non-generic generalisations generics quantify over in §4.3 of this chapter.

[69] For reasons set out in Chapter 1, I am assuming that generics contain a covert quantifier, Gen, which I am arguing is a restricted existential quantifier that ranges over non-generic generalisations. However, the idea that generics quantify existentially is also compatible with believing that some pragmatic mechanism gives generics this quantificational force albeit less straightforwardly.

Non-specific generalisations still generalise about members of a kind by linking it to some property, but such generalisations are less specific about the nature of that link. Specificity, including in generalisations, comes in degrees. Some generalisations contain a lot of information about the nature of the link they ascribe, some very little, and most are somewhere in between.

(2) 87% of Scots are vitamin D deficient

(3) Most Scots are vitamin D deficient

(4) Many Scots are vitamin D deficient

For example, (2) is a generalisation about Scottish people and vitamin D deficiency. The generalisation is more specific than (3), which in turn is more specific than (4). 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 non-generic generalisations.

Generics are non-specific in this sense because they quantify existentially over more specific generalisations. The status quo in the generics literature is that even though generics don't contain quantifier expressions at the surface level, they do contain a covert quantifier called 'Gen', a matter I examined in Chapter 1. Historically, the dominant view in the literature has been the idea that Gen is a restricted universal quantifier which quantifies over normal individuals or normal worlds.[70]

On such a view, generics of the form *Fs are G* say something like that all normal Fs are G, or that all Fs are such that the most normal possible world(s) for them are ones in which they are G.[71] Instead of thinking that Gen has universal quantificational force of some kind, I want to explore the idea that its quantificational force is existential.

2.2 Existential Quantification

If generics do quantify existentially, why think that the domain of quantification contains non-generic generalisations and not simply members of the kind the generalisation is about? Generics of the form *Fs are G* don't simply say that some Fs are G. Otherwise the following generics would be true.

(5) Humans eat stones

(6) Sheep eat pizza

While some humans have eaten stones and some sheep have eaten pizza, we wouldn't want to say that either (5) or (6) are true. Instead, the existential quantification must be over entities other than individual humans or sheep. My proposal is that the entities generics quantify over existentially are themselves

[70] These are the normalcy accounts I discussed briefly in §2 of Chapter 2. I discuss a refined version of them in §4 of this chapter.
[71] See Asher and Morreau (1995) and Asher and Pelletier (2012) for influential articulations of such views.

generalisations, albeit ones that are more specific than generics. I discuss which such generalisations are quantified over in the sections to follow but want to first give an account of how this existential quantification arises.

3 Existential closure

In Chapter 1, I mentioned that same sentences that can receive generic interpretations in English, bare plurals, definite, and indefinite singulars, can also receive existential interpretations.

(7) Sheep are in the garden

(8) The sheep is eating a lollypop

(9) A sheep stood in the road

(7)-(9) don't receive generic interpretations. They are about particular sheep, not sheep in general. While accounts explaining this phenomenon differ, many of them make use of the notion of 'existential closure' to explain how sentences like (7)-(9) receive existential readings.[72] Existential Closure occurs when variables that are otherwise free are existentially bound. For example, Irene Heim (1982) proposes an analysis of indefinites, according to which they have no quantificational force on their own but contribute a free variable. This free variable can be bound by overt quantifiers, but, when this doesn't happen, it is bound existentially instead.

For example:

(10) A cup of tea always gets cold too quickly

(11) A cup of tea is on the table

In both (10) and (11), the indefinite noun phrase *A cup of tea* introduces a free variable. In (10) it is bound by *always,* a quantifier.[73] However, in (11), there is no such quantifier to bind the variable, which is why it is bound existentially. Some who want to explain why sentences of the same syntactic form, in English bare plurals, indefinite and definite singulars, can receive either generic or existential readings also make use of the notion of existential closure. For example, Krifka (2003) argues that bare plurals get an existential reading when the variables they contain are not bound by a generic operator. Krifka thinks of the generic operator as a restricted universal quantifier, quantifying over normal individuals. However, if it is instead possible to think of GEN as a restricted existential quantifier, then this raises the question of whether quantification in generics itself could not be a type of existential closure.

My idea is that the existential quantification that I've argued arises in generics could be thought of as arising as the result of existential closure. When a generalisation does not contain a quantifier expression,

---

[72] For instance, see work by Kamp (2002) and Heim (1982) and Diesing (1992).

[73] Or, speaking more carefully and in keeping with my use of 'quantifier' in Chapter 1, the free variable introduced by the infinite NP is bound by the quantifier that 'always' refers to.

a free variable introduced by the NP is instead bound existentially in one of several ways. This could also explain why BP, IS and DS sentences do not receive generic interpretations when they contain overt quantifier expressions, as is the case in (10). If the free variable is not bound by another quantifier, the sentence receives an existential reading, though whether the existential quantification is over individuals, events, relations, or, in the case of generics, other generalisations is dependent on syntactic and pragmatic features.[74]

The existential closure proposal could also help develop the definition of generics offered in chapter 1. There, I argued that generics should be defined as lacking quantifier expressions, but noted that this definition would have to specify where in the generic there cannot be any quantifier expressions. The existential closure proposal could help fill in this gap. The idea would be that quantifier expressions can occur in generics but not in such a way so as to bind the free variable introduced by the NP. This free variable is subsequently bound existentially by way of existential closure. [75]

The existential closure component is not an integral part of the generality account. Existential quantification could arise for other reasons after all. However, if plausible, the existential closure hypothesis could explain why Gen usually isn't marked in natural languages. It's because existential closure kicks in precisely when there is no overt quantifier present in a generalisation that it receives a generic reading. Having characterised the source of existential quantification in generics as a type of

---

[74] An anonymous reviewer for Synthese suggests that those who think of non-specificity in generics as a by-product of metasemantic phenomena like contextualism or semantic incompleteness, the 'metasemanticists', could also argue that generics uttered in indeterminate contexts express weak propositions as a result of existential closure, without needing to posit non-specificity separately. My best guess of what such an account would look like is as follows: Where generics are uttered in indeterminate contexts, whatever values the generic token is compatible with in the context would be existentially quantified over. This explanation would make the same predictions about the propositional content of such generic tokens but would not need to posit non-specificity independently. 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.
Here are four considerations that move me towards the semantic non-specificity view:
-My proposal can explain how existential closure is triggered by the absence of a quantifier expression at the surface level of sentences in keeping with most other theories that invoke existential closure, whereas the alternative explanation must think of it as a pragmatic process.
-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.
-I will suggest that context-sensitivity may not be needed to account for truthmaker and resulting use variability of the kind that motivates contextualists, in which case non-specificity and whatever metasemantic phenomenon is posited may be compatible, but the latter may no longer be required to explain variability.
-In the next three chapters, I argue that having non-specific generalisations be both the semantic content of generalisations and the contents of propositional attitudes, especially belief, is theoretically useful in various ways that are difficult for the metasemanticist to accommodate.

[75] The difference between the proposal of this kind, and the proposal described in the previous footnote lies in whether existential closure occurs only when (context-sensitive) generics are uttered in indeterminate contexts, or whether it occurs by default due to an absence of suitable quantifier expressions.

existential closure, I will now move on to develop the account I sketched earlier in more detail by comparing it to its closest relative in the generics literature: Nickel's (2016) normalcy account.[76]

4 Comparison with Nickel

4.1 Nickel's account

Bernhard Nickel (2016) presents an account of generics that makes use of existential quantification. Nickel's starting point is to modify a relatively standard normalcy account.

(12) Dogs bark

(13) All dogs bark

If (12) says something about what all normal dogs are like it can, in contrast to unrestricted universally quantified generalisations like (13), tolerate dogs which don't bark, as long as they are in some sense abnormal. However, Nickel departs from the standard normalcy approach by acknowledging that there can be more than one way of being normal for a given kind and property. For him, the truth conditions of generics are as follows:[77]

> *Fs are G* is true iff there is a way of being a normal F with respect to G-ness, such that all those Fs which are normal in this way are G. [78]

For Nickel, the content of G determines the relevant respect of being normal. For example, if it is *x barks,* then the relevant normality is normality with respect to characteristic noises. The generality account is structurally similar to Nickels in that, according to both, generics quantify existentially. Where they differ, is with regards to the domain of quantification. For Nickel, generics quantify over ways of being normal, whereas I argue that they quantify over non-generic generalisations. Consider the following generic to see why Nickel acknowledges that there can be more than one way of being normal:

(14) Sea turtles are long-lived

Nickel can explain how it is that tokens of (14) can be true despite the fact that the vast majority of sea-turtles die shortly after birth. While it is not statistically normal for sea-turtles to be long lived, there is another way in which it is normal for them to do so. To see how, we need to examine Nickel's notion of normality.

---

[76] The only other similar account I am aware of is given by Carlson (2008), who argues that generics quantify existentially over inductive patterns. Carlson's account is relatively schematic and hasn't received as much attention as Nickel's. Nevertheless, many of the things I identify as advantages of Nickel's account, and want to make use of in my own, also apply to Carlson's account.
[77] Nickel (2016) goes through a range of formulations to deal with various complications in Chapter 3. See p.68 for his final formulation.
[78] The 'with respect to G-ness' proviso is introduced to deal with the fact that members can be normal with respect to one characteristic while being abnormal with respect to another. I ignore it for the rest of my discussion.

4.2 Nickel's notion of normality

Nickel's notion of normality in the truth-conditions of generics is spelled out by a theory of genericity, which is connected to his semantic theory for generics through a bridge principle.[79] For Nickel, generics are about what's characteristic for members of a kind.[80] Being characteristic is defined relative to a set of explanatory strategies selected by the context in which the generic in question is uttered. Each way of being normal is therefore underwritten by a causal mechanism that is recognised by contextually salient explanatory strategies.[81]

For example, if (14) is uttered in a context in which explanatory strategies from evolutionary biology are salient, then being long-lived can correctly be ascribed as being characteristic of sea turtles, because there is at least one way of being long-lived for sea-turtles that is recognised by the contextually salient explanatory strategy, namely that sea turtles, once they survive infancy, tend to be long-lived because of features of their biological make-up that they have as result of evolutionary processes. Nickel's account of normalcy in terms of mechanism is nuanced and complicated so I can only give a brief overview into the reasons that move me to look for a structurally similar account that eschews normality.

First, Nickel assumes that mechanisms are causal. This leaves unclear how his account applies to generics with arguably causally inert subject matters, such as in mathematics, theology, or ethics.

(15) Numbers are in platonic heaven

(16) Souls are immortal

(17) Ethical properties are irreducible to non-moral properties

Many believe that numbers, souls, or ethical properties are not the kind of things that can causally interact with the world. For them, generics about such subject matter couldn't, on Nickel's account, be true because they couldn't be related to the relevant properties via causal mechanisms. So at least for them, either (15) - (17) or Nickel's account must be false.

Second, even disregarding the possibility of domains with causally inert subject matter, Nickel's account requires that the generics we want to count as true are underwritten by mechanisms. This is problematic because our talk about the world doesn't neatly coincide with causal mechanisms that cause generalisations to be true. I will illustrate this general point using a more local problem identified by Alnica Visser.[82]

Visser focusses on generics about anthropic kinds. Anthropic kinds are classes of individuals grouped together for human purposes and include things like *vegetables* or *wood*. Many of the generics we use in

[79] Nickel (2016:197)
[80] Nickel (2016:181)
[81] Nickel (2016:198)
[82] Visser's paper hasn't been published. See Bibliography for details.

everyday speech are about exactly such categories, including many that initially seem to be about scientific kinds. The issue is that individuals grouped using anthropic labels frequently do not share in a common mechanism that explains how they are related to some property. For example, one of the most frequently used generics by Nickel is.

(18) Ravens are black

As Visser points out, ravens do not form a unified class within the explanatory strategies, i.e., zoology, employed in standard contexts. Different birds we call 'ravens' are no more closely related to each other than they are to other birds. Hence, while there are various evolutionary mechanisms that would be suitable for explaining how having black feathers can be characteristic of birds, the kind in question, ravens, isn't recognised in the relevant domain of inquiry. However, this also means that there cannot be a suitable mechanism, linking the two.

Nickel recognises this issue, though not with regards to ravens. He says:

> 'Where explanatory strategies are associated with a systematic arena of inquiry regulated by the world, the kinds whose characteristic properties are underwritten by these explanatory strategies are the kinds recognized by that discipline, and only those kinds' (2016:209).

To deal with cases in which seemingly true generics are predicated of kinds not recognised by the salient explanatory strategies, Nickel introduces the notion of a derivative generic. These are generics that predicate a property of what he calls a 'kind-derived subgroup.' (2016:206) According to Nickel, a generic concerning a kind-derived subgroup can be true even though the kind may not be recognised in a particular domain of inquiry, as long as there is a mechanism that links the members of the subgroup to the property in question. So, as long as there is one mechanism that links members of a subgroup, in this case those birds we call 'ravens' to the property of having black feathers, the property is characteristic relative to explanatory strategies that recognise this mechanism.

However, as Visser points out, the issue isn't merely that certain kinds, like ravens, are not recognised in the relevant domain of inquiry. Rather, many members of anthropic kinds who have similar properties, say concerning the colours of their feathers, have those properties in virtue of differing causal mechanisms because they are only grouped together for reasons of human convenience. Consequently, there is no *one* relevant mechanism that links ravens and the property of having black feathers.

This might initially seem like a pedantic issue with Nickel's articulation of his theory, but the issue is a more general one. Nickel wants to explain the semantics of generics in a way that anchors their truth in causal mechanisms that are recognised by distinct explanatory strategies. However, as Visser's example suggests, there can be mismatches between our explanatory strategies, the kinds we talk about and the properties we ascribe to them. Our explanatory strategies can be ill-suited to the domain to which they are applied, but this doesn't appear to make the relevant generics false, as in the case of (18). Is it really reasonable to think that our every-day use of generics can only be sanctioned by explanatory strategies

that are responsive only to specific, causal mechanisms? The types of non-specific uses of generics, like Aoife's utterance, that motivated my search for an alternative view also seem to suggest that many contexts in which generics are uttered are not ones in which only *one* discrete explanatory strategy is salient.

However, as Nickel's account of normality is given in terms of mechanisms, Nickel cannot simply drop or dilute the mechanism requirement without making his account uninformative. The types of claims generics can be used to make already exhibit such great variety that a very flexible notion of normality is required to accommodate them all. Nickel's acknowledgement of different ways of being normal is a symptom of this. If there were nothing to draw together what different ways of being normal involved, then there seems to me to be little explanatory value in appealing to normality in the first place.

In contrast, the generality account makes no use of the notion of normality. It does not understand generics as really being about what's normal for members of a kind. Some of the non-generic generalisations that generics quantify over will be normality-flavoured ones, but there is no requirement that they all need to be. According to the generality account, generics are not to be understood as sentences that are about *one type* of generalisation as opposed to another: instead, they are about all of them.

4.3 The domain of quantification

If no more general notion acts as a criterion for *which* non-generic generalisations generics quantify over, then should we think that any generalisations at all are eligible to make generics true? I don't think so. For example, if most horses liked wearing horseshoes, but none ever in any sense *did*, this doesn't seem to me to suffice to make (1), *Horses wear horseshoes*, true. So not just any generalisation whatsoever about horses and the wearing of horseshoes can be in the set of non-generic generalisations that can make (1) true. Instead, the domain of non-generic generalisations needs to be restricted.

Exactly which non-generic generalisations can be quantified over is not a matter I can or want to specify at this point. Doing so would require empirical work that hasn't, as far as I am aware, been undertaken. What is needed would be experimental testing of which non-generic generalisations speakers take to make true the corresponding generics. These seem likely to me to include generalisations corresponding to many of the types of generalisations in terms of which generics have been analysed: claims over 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 the types of generalisations that can by requiring that they all share some underlying feature in the way Nickel does, especially when the notion used for this purpose is either too restrictive or uninformative.

One thing that is important to point out that each of the generalisations generics quantify over existentially can themselves be non-specific to varying extents. For example, there are different kinds of statistical generalisations, such as absolute and comparative prevalence generalisations, each of which can themselves also be vague. 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. Having sketched the basic generality account in some detail, I want to note some of its advantages before anticipating an objection.

5 Advantages

Two advantages are, first, that it can explain non-specificity and, second, that it has the logical structure required to account for the interaction between generics and logical connectives.

5.1 Non-specificity

I argued in the previous chapter that it is important to be able to account for non-specific uses of generics. According to the generality account in its basic guise, a generic token of the form *Fs are G* simply says that at least one salient non-generic generalisation about being an F and G-ness is true. This approach can explain non-specificity as a feature of generics simply in virtue of the type of propositions generics express semantically. Given that generics existentially quantify over non-generic generalisations, any one of those generalisations being true makes the relevant generic true. Because the propositions generics express are themselves non-specific, speakers need not have any more specific intentions to successfully communicate them. Because the generality account locates non-specificity in the semantic content of generics, no additional explanation of their non-specificity is needed.

5.2 Strong negation

The second advantage of the generality account that I want to draw attention to is that it can account for how generics interact with negation and other logical connectives.

(19) Rabbits eats grass

(20) Rabbits don't eat meat

(21) Rabbits eat meat

I will call sentences like (20) 'negated generics'. Carlson (2008) notes that negated generics seem to be especially strong. For example, while (19) seems to allow for plenty of rabbits that don't eat grass, (20) doesn't allow for many, if any, meat-eating rabbits. This claim must be qualified. While some negated generics are strong in the sense Carlson describes, other aren't.

(22) Pandas don't live in Europe

(23) Doctors don't wash their hands

(24) Librarians don't go clubbing

(22) - (24) seem like they might be true even though quite a few pandas live in Europe, a lack of hand-washing among doctors is a well-recognised problem, and some librarians do go clubbing. So, it seems that while some negated generics are notably strong, others aren't.

The generality account has the logical structure to be able to explain this phenomenon as a simple case of scope ambiguity. If negation takes wide scope over the existential quantifier, we get a stronger reading than when it takes narrow scope. (20) is therefore the wide scope reading of the negation of (21). This explanation can account for the varying strength of negated generics. On a wide-scope reading, they are true iff *none* of a set of non-generic generalisations concerning the kind and property in question are true, whereas on a narrow scope reading, they are true iff at least one of several non-generic generalisation is not.[83]

If my explanation so far is right, then why do we still sometimes take sentences like (20) to be the unique negation of sentences like (21)? For example, if someone were to utter (21) to express their mistaken belief that rabbits eat meat, someone else might respond by uttering (20), seemingly to disagree with them. We can explain this behaviour as an instance of *Neg-raising*. Horn (2020) defines neg-raising as 'the strong tendency in many languages to attract to the main verb a negative which should logically belong to the dependent nexus [=clause]' (2020:199). A well-known example of this is when we equate a lack of belief that P with a belief that not-P, such as a lack of belief in the existence of ghosts with a belief that ghosts don't exist. In any case, the tendency to get confused about negations doesn't seem as though it needs to be accounted for specifically by a theory of generics but can be explained as part of a wider theory of how we understand negated claims.

Although I have focussed on the interaction between generics and negation in this subsection, this is only one interaction between generics and linguistic phenomena that the generality account is well-equipped to explain. Nickel has given detailed accounts of the interaction between generics and various logical connectives, specifically disjunction and conjunction.[84] In doing so, he makes use of the logical form of generics as featuring a restricted existential quantifier.

(25) Elephants live in Africa and Asia

a) Many elephants live in Africa

b) Many elephants live in Asia

Standard normalcy accounts struggle with sentences like (25). They understand generics as something akin to claims about all normal elephants. However, it definitely isn't true of all normal elephants that they live in Asia *and* Africa. Nevertheless, (25) seems true. Nickel can explain the truth of (25) in terms of

[83] Carlson (2008) pursues a similar line of argument.

[84] Nickel (2016) dedicates Chapters 5 and 6 to examining these matters.

generics quantifying over ways of being normal for elephants. One way of being normal for elephants is to live in Africa, another to live in Asia.[85] Because the generality account is structurally similar to Nickel's account and only differs in the domain of quantification, it can reap the same benefits. Because, as Nickel puts it, there is one way of being normal for elephants such that they live in Africa, but also another such that they live in Asia, there are also corresponding true generalisations about each elephant demographic, for example as in (25a) and (25b).[86] Consequently, the generality account can explain the truth of (25) by making use of Nickel's work.

I hope that the above has given some initial reasons to be interested in the generality account. In the next few chapters, I will be drawing attention to various additional explanatory benefits of the generality account that hinge on its ability to provide the tools for an account of non-specific generic beliefs. However, there is an obvious objection to the account as I have presented it so far that I must address first: that it overgenerates true generics.

6 Overgeneration

6.1 The problem

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 seem to be many cases in which a seemingly salient non-generic generalisation about a kind and property is true, but the corresponding generic is false.

For example, consider the widely discussed case of:

(26) Books are paperbacks

Even though most books are paperbacks, utterances of (26) still sound strange and, to some, straightforwardly false.[87] However, according to the generality account, (26) should be true because most books are paperbacks.

The problem isn't just limited to sentences like (26). 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 different contexts of utterance. But in its basic version, the non-specificity account predicts that the truth-conditions of generics are stable in that they existentially quantify over the

[85] Nickel (2016:128)

[86] In this case I am assuming that (25a) and (25b) are non-generic generalisations that would be judged as apt to make true generics. If this turned out not to be the case, I think it is plausible that other generalisations would take their place, for example ones about it being natural for elephants to live in either place.

[87] Leslie (2008)

same set of non-generic generalisations and would consequently be true or false in all contexts. So, don't generic tokens express different generalisations in different contexts after all?

6.1.1 Contextualism to the rescue?

One straightforward way of warding off both these worries is to combine non-specificity 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. According to a contextualist generality account, non-specificity will be a feature of generic types, and to varying extents, generic tokens.

The features of contexts that determine the semantic values of generic tokens could be the speaker's intentions[88] or propositional attitudes[89] shared by the conversation's participants. Relevant propositional attitudes would also reflect background beliefs that participants hold. As a result, some generic tokens would, at least in standard contexts, quantify only over a limited set of specific generalisations. For example, given that most participants in conversations know that there are flightless birds, utterances of *Birds fly* wouldn't usually quantify over the claim that all birds fly because the shared background beliefs of participants are incompatible with this universal generalisation.

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 the members of the kind and the property in question is true, but say little about the nature of that generalisation, for instance whether it ascribes absolute or comparative prevalence, a dispositional link, characteristics, or even a normative connection. 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 would alleviate some overgeneration worries, at least initially. Many of the generic tokens that the invariant version of the generality account would controversially predict to be true, could be false on the contextualist version simply because the non-generic generalisation that is true would not be contextually salient. We could also account for use variability in the standard way contextualists do. A generic type can be used to convey a generalisation of a particular flavour or quantificational force in one context. Yet in another the context, the 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. In this way, the proponent of the

---

[88] King (2013)

[89] See Stalnaker (1978) and Heim (1982). I include in these beliefs about the investigative purposes of the conversation which determine what has been termed the question under discussion (QUD) after Roberts (2012).

contextualist generality account can easily account for use-variability without having to resort to propositional nihilism or pluralism to explain non-specificity in some generics.

6.2 Embracing true generics

The above is a line of argument against the overgeneration worry that is open to the proponent of the contextualist generality account only. However, I don't think that these considerations are decisive for several reasons: For now, I will limit myself to only discussing overgeneration as a positive reason for a contextualist version and argue that it can be responded to. However, I will, at various points in the next three chapters, return to considering whether a contextualist version of the generality account is to be preferred to an invariant version. For now, I don't think that overgeneration worries provide good reasons to accept a contextualist version for three reasons: the relief they provide 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.

6.2.1 Everybody's problem

First, the overgeneration problems simply resurface for the contextualist, or indeed those in favour of the semantic incompleteness view, as questions about why a given non-generic generalisation cannot be salient in a given context or cannot be completed in a certain way. Why can't we utter (26) intending to make the claim that most books are paperbacks and thereby convey it?

Even accounts that do not posit context-sensitivity or semantic incompleteness struggle with (26) and other sentences like it. For example, on Leslie's account, (26) is what she calls a 'majority generic', the type of generic that is made true by the fact that most members of the kind have the property in question. However, the reason for why (26) is not true, according to Leslie, is that she requires, in addition, that any counter-instances, i.e., any members of the kind that don't have the property are 'negative' and not 'positive' counter-instances. What she means by this is complicated, but, roughly, the idea is that it is significant that those books that are not paperbacks don't just *lack* the property of being a paperback, but instead have the positive property of being a hardback book.[90] However, as Nickel (2016) and Sterken (2015b) discuss, this distinction comes under pressure in a variety of ways, and ultimately leads Leslie's account to serious under-generation problems. Questions about sentences like (26) affect everyone writing on generics in one form or another, whether as semantic or meta-semantic worries, and so while an answer to them should be sought, they shouldn't count as a decisive point against the generality account.

[90] Leslie discusses the distinction in (2008:33-34)

6.2.2 Truth-value judgements

The second thing to say is that the presentation of truth-value judgements in the generics literature is sometimes overly simplistic in a way that catastrophises overgeneration worries. The verdicts on which those worries are based, i.e., that (26) is false, are often assumed to be obvious and thus little further evidence is provided.

Even within the generics literature, there is disagreement about the truth-conditions of some generics. For example, Krifka et al (1989) 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) gives reason to think this can be disputed.[91] 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. With regards to (26), Leslie (2008) assumes that it is false, whereas Nickel (2018) doesn't.

Further, the studies we do have show a complicated picture with regards to people's truth-value judgements. For example, Prasada et al (2013) conducted an experiment in which participants were given several generics and asked to judge their truth-value. Participants were asked to evaluate generic sentences along a 7-point scale with -3 labelled 'definitely false' and +3 labelled 'definitely true'.

One interesting finding concerns the truth scores participants gave for supposedly false generics. The highest rating given was a mean judgement of 2.91 for *Airplanes have wings*, the lowest was a mean of -0.91 for *Restaurants are Chinese restaurants*. The latter is quite interesting because it shows that even among a list of supposedly false generics, the generic that was judged to be the falsest was still judged to be quite a long way away from being definitely false (-3). Several of the generics that were given as examples of false generics were in fact judged by participants as more true than false. For instance, the mean for (26), *Books are paperbacks*, was 0.21 and another widely discussed example, *Bees are worker bees*, received a mean judgement of .96. These sentences did receive lower truth-value judgements that many of the other generics but were still judged to be more true than false. Of course, this is only one piece of research and the truth value judgements may be explainable by other means.[92] However, I do think that they give us reason not to assume that the example sentences that are often taken as clear examples of false generics are actually false.

Further, the relevant studies do not distinguish between whether the generics in question are false, or just sound bad or otherwise odd. This also leaves open the possibility that we agree that something is wrong with sentences like (26), but that that need not be the fact that they are false.[93] For several reasons then,

---

[91] Wasserman (2011:432)

[92] I use this study because it supplies a break-down for individual sentences. Other research gives mean acceptance rates or truth-value judgements for the whole category of supposedly false generics which can be quite diverse. See for example Leslie et al (2011) and Lazaridou-Chatzigoga et al (2019). The latter do mention (26) specifically and conclude that 'people generally find judging these kinds of statements difficult' (2019:635).

[93] For example, Nickel (2018) argues that sentences like (26) are not false, but uninterpretable because of constraints brought about by the investigative purposes of the contexts in which they are standardly asserted.

more evidence is required to substantiate the idea that many of 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 (26). But what about the ability to explain the observations of what looks like truth-conditional token variability?

7 Pragmatic contents

It will by now have become clear why I insisted on talking about truthmaker variability, not truth-conditional variability in Chapter 2. Truthmaker variability is entirely consistent with the invariant version of the generality account as any of a list of non-generic generalisations can make true the relevant generic. However, truth-conditional variability isn't. According to the generality account in its invariant form, the truth-conditions of generics are non-specific but stable. So, what to make of the examples of what looked like truth-conditional variability that Sterken and Nguyen used to motivate their theories of generics as context-sensitive or semantically incomplete?

I will argue in this section 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. 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 use are not, on the picture I propose, instances of truth-conditional variability, but of variability in use based on the pragmatic effects of generic tokens. In making my argument, I will draw on Haslanger's work on generic implicatures.

7.1 Haslanger

Sally Haslanger (2011, 2014) argues that assertions of generics convey distinctive implicatures in addition to their semantic content.[94] According to her, assertions of most generics of the form *Ks are F* convey additional information about the nature of the connection between Fs and G-ness by means of a conversational implicature.[95] [96]

---

[94] Haslanger offers an initial account in Haslanger (2011) and refines it in (2014). I will focus on the 2014 account for reasons given by Saul (2017) against the (2011) account.

[95] Haslanger argues that generics that have a clear statistical flavour (In Leslie's terms, generics that are clear majority generics) do not convey the implicatures in question.(2011:30) Because I believe that generic sentences are non-specific, I don't think that any generic sentences are clear majority generics. Consequently, I think that an account of whether utterances of generics generate essentialist, normative, or other implicatures should be more general in its scope and apply to all generics. I don't focus on this issue with Haslanger's proposal here but note that it does not apply to my proposed semantic extension of her account.

[96] Haslanger leaves open whether the content might be presupposed instead of implicated:
'I'm actually not sure whether it is better to consider it an implication or a presupposition. I'm willing to adjust my account to accommodate evidence for either' (2014: 370 Fn 9).
I think she is right to focus on implicatures and not presuppositions. Presupposed content, in contrast with content implicated conversationally, projects when embedded and is not concealable, whereas the pragmatic content Haslanger hypothesises is

Haslanger (2014) gives a two-step answer as to how these implicatures are generated. First, assertions of most generics of the form *Ks are F* implicate that 'the connection between the Ks and F holds primarily by virtue of some important facts about the Ks as such, or by virtue of what it is to be a K' (2014:370).

(27) Fire fighters are courageous.

For example, an assertion of (27) implicates that there is a robust connection between being a fire fighter and being courageous. A fire fighter's being courageous doesn't just occur accidentally, but as a result of some important fact about them. In addition to the robustness implicature, Haslanger argues that additional assumptions that are part of the common ground, in particular an essentialist assumption[97] and a normative assumption[98], generate stronger, explanatory implicatures.[99] If these assumptions are operative in the contexts in which generics are uttered, implicatures to the effect that members of the kind have the property in question as a result of their generic essence and that this ought to be the case are generated. For example, if the common ground of a conversation in which (24) is asserted contains an essentialist assumption, then an additional implicature, namely that the important fact in virtue of which fire fighters are courageous is their essence, is generated. If the common ground also contains a normative assumption, a further implicature, namely that fire fighters ought to be courageous, is generated.

Though I agree with Haslanger's general approach, I disagree with its implementation. Without further explaining the content of the robustness implicature Haslanger, it is unclear what it means to say that their courageousness is due to some (which?) important facts about them. Instead, Haslanger should drop the robustness implicature and simply adopt a suitably weak semantics of generics, such as the generality account. This locates something close to the content of the would-be robustness implicature in the semantics of generics. Further, the semantic content can, in the way Haslanger explains, be combined with cognitive biases or other background beliefs to explain how stronger implicatures are generated.

Whether implicatures are generated depends, according to Haslanger on whether or not essentialist and normative assumptions are present in the common ground of a conversation, which in turns depends on whether the context 'calls for an explanation' (2014:16). This part of Haslanger's proposal is very schematic and could, I propose, benefit from reflections on the types of inferences generic beliefs dispose

---

cancellable (c) and does appear to project (a and b).
(27) Fire fighters are courageous
a) Are fire fighters courageous?
b) Fire fighters are not courageous.
c) Fire fighters are courageous, as a result of their circumstances, and not their natures.
For the same reasons, the implicatures must be conversational and not conventional.

[97] 'Essentialist Assumption: Robust (meaningful?) regularities are not accidental. They are due to the natures of things' (2014:13).

[98] 'Normative Assumption: Things should express their natures and under normal circumstances they will. Abnormal circumstances are not good and should be avoided or changed' (2014:13).

[99] For a detailed discussion of both the robust and explanatory implicatures, see Rosola (2019).

us to. I will discuss the nature of generic beliefs in more detail in the next chapter, however, here is a quick preview.

We can use generics to communicate non-specific generalisations which interact with our interlocutors' prior beliefs to induce more specific beliefs. This is a general feature of non-specific language. For example, imagine I tell you that I went to the bakery yesterday. If you know that there are only two bakeries, I have access to, you will likely believe that I went to one of the two. If it turns out I went to neither you will be sceptical of my claim and might well suspect that it was false.

(28) Ducks lay eggs

Likewise, while (28) itself only makes the claim that *one* of several non-generic generalisations about ducks and the laying of eggs is true, our background beliefs combine with the belief we form on the basis of hearing (28). Assuming that most of us have the background belief that it's female ducks that bear young and that less than half the population of ducks will be made up of fertile female ducks, we can restrict which generalisations about ducks laying eggs are those that are likely to make (28) true. So, while an utterance of (28) semantically may only express the weak, non-specific generalisation, it can, in a context with suitable background assumptions, convey additional conversational implicatures about which restricted set of non-generic candidates are the possible truthmaker for the generic. In contrast to contextualist versions of the generality account, this restriction is pragmatic, not semantic.

Hence, even though (28) doesn't have a strong semantic meaning, it, in combination with our background beliefs, conveys more specific content. In cases in which an utterance of a generic pragmatically conveys one or several more specific generalisations, one of which then turns out to be false, we may think that the utterance of the generic sounds bad (it made us believe something false after all) even though it is, strictly speaking, true.

This explanation also has the potential to explain the ways in which generics can be used to convey different, non-generic generalisations in different contexts of utterance, even though the generic tokens themselves have stable semantic contents.

(29) Nurses are caring

While (29) has the same semantic content across different contexts of utterance, it interacts with cognitive biases, background beliefs as well as general relevance expectations to convey more specific generalisations via implicatures. For example, in contexts which contain normative assumptions, an utterance of (29) will give rise to the more specific implicature that nurses should be caring. Conversely, in contexts that contain an essentialist assumption, utterances of (29) generate the implicature that nurses are caring in virtue of some shared nature. This means that (29) can be used to convey more specific generalisations about the link between being a nurse and being caring, even if its semantic value is less specific.

Other implicatures generated by generic tokens depend less on normative or essentialist biases, and more on general relevance expectations and their interactions with beliefs about the relevant kinds or properties.[100] The context in which generics are uttered can determine whether the generic uttered is relevant or not. When a context is such that only a proper subset of non-generic generalisation 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. For example, a context in which two evolutionary biologists are discussing horses and the various adaptations horses have undertaken in response to its environment, uttering (1), *Horses wear horseshoes,* is 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.

Because of this interaction between investigative aims of a context and relevance expectations, speakers can use generics to pragmatically convey more specific generalisations. For example, the evolutionary biologist who utters (1) in the course of their TV programme knows that their audience will reasonably interpret their utterance as being about what horses are like as a result of evolutionary adaptations, or some other, similar non-generic generalisation, given that nothing else would be relevant in the context of utterance. In contrast, in a conversation between horse breeders about horse-care related budgets, someone who utters (1) would convey relevant implicatures about the prevalence of horseshoe-wearing among horses and likely none about how horses have evolved.

I discuss the cognitive background that gives rise to these implicatures in more detail in Chapter 4 and a reasons for why speakers may want to use generics instead of overtly quantified generalisations in Chapter 5. However, at least for now, the theoretical avenues available to us suggest that we can use pragmatic means to explain use variability in a way that doesn't require contextualism.

8 Conclusion

In this chapter I have presented a new account of generics. The generality account says that generics feature a restricted existential quantifier, Gen. Gen quantifies over sets containing non-generic generalisations. The generality account does a good job of accounting for non-specificity in virtue of existential quantification over non-generic generalisations and can explain the variety of ways in which generics interact with logical connectives. 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.

The proposal I have presented in this chapter so far has been that, instead of thinking of non-specificity as a by-product of a phenomenon like context-sensitivity or semantic incompleteness, we should think of

[100] For other recent 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.

it as a feature of the semantic content of all generic sentences. In the next three chapters, I will illustrate how the non-specific generalisations are useful explanatory posits, as semantic values but also as *objects of belief*. This extension, I hope, will add to the argument made in Chapter 2 to build an abductive case for non-specificity approaches, and the generality account in particular.

Chapter 4: Generic Belief

1 Introduction

The vast majority of work on genericity has focused on generics as linguistic items. Not many theorists have focussed on genericity as a feature of mental states. One exception to this rule is Sarah-Jane Leslie, who argues that generics give voice to cognitively basic generalisations.[101] In this chapter, I evaluate Leslie's approach and propose an alternative. The chapter is divided into two parts.

In the first, I evaluate Leslie's framework of generics as cognitively primitive generalisations and take issue with two aspects of her account: one empirical and one theoretical. First, I argue that the currently available empirical evidence is not as conclusive as Leslie presents it as being. Second, Leslie's distinction between the conditions under which a generic is true, and those under which the generalisation it gives voice to is accurate, leads to a strange disjointedness in her account that limits its explanatory value.

In the second, I offer an alternative. I develop an account of generic belief by extending the generality account I presented in the previous chapter. Just as non-specific generalisations can be expressed using generic sentences, they can also be the objects of propositional attitudes like belief. On the account I propose, generic beliefs are distinctive attitudes towards non-specific generalisations. Such beliefs may co-exist alongside other beliefs in more specific generalisations, or they may stand alone. I argue that this framework provides a more cohesive account of genericity in mental states and can explain some empirical data concerning the inference patterns generic beliefs give rise to.

2 Generics as default hypothesis

2.1 Leslie's account

Sarah-Jane Leslie (2007, 2008, 2012) has built her account of generics around the idea that they express distinctive mental states. According to Leslie, these mental states are cognitively basic, 'default' (2012:31) generalisations, which generics give voice to. Leslie explicitly aligns herself with a dual-process approach, according to which our reasoning processes divide into two systems.[102] System 1 processes are fast, automatic, and unconscious, whereas System 2 processes are slower, controlled, and conscious. Leslie argues that generics express generalisations made by System 1 processes. She describes these generalisations as 'cognitively fundamental' and 'primitive' (2012:31) and contrasts them with sentences expressing quantificational generalisations, which according to Leslie, are the result of System 2 processes.

[101] Some other exceptions are Cohen (2004) and Prasada (2000).

[102] Leslie uses Kahneman's framework as set out in Kahneman (2002, 2011) and Kahneman and Frederick (2002). See Frankish (2010) for an overview of dual process theories and how Kahneman's work relates to that of others within the tradition.

2.2 Empirical predictions

As Leslie's account is embedded within a dual systems approach, it makes several empirical predictions.

First, according to Leslie, children can make the cognitively primitive generalisations that are expressed by generics before they can make cognitively more sophisticated System 2 generalisations. This should be demonstrated by the ages at which children acquire generics in comparison to the age at which they acquire overtly quantified generalisations.[103]

Second, if generics express cognitively fundamental generalisations, then we can expect that, in contexts in which cognitive processing demands are high, people will inappropriately make generic generalisations instead of overtly quantified ones.[104] This in instance of a wider phenomenon Daniel Kahneman (2002) describes whereby the judgements made by Systems 1 and 2 processes sometimes come into conflict and produce System 1 judgments that have to be consciously overridden by Systems 2 processes.

In one famous example, participants are told that a bat and a ball cost $1.10 and that the bat costs $1 more than the ball.[105] When asked how much the ball costs, most answer that is costs ten cents even though the answer is five cents. Kahneman argues that our intuitive and automatic System 1 processes give us an answer of ten cents and have to be consciously overridden by System 2 processes to arrive at the correct conclusion.

Leslie cites a range of experiments which she takes to bear out these predictions. I will give brief descriptions of some of the main pieces of evidence Leslie takes to be in favour of her view.

2.3 Acquisition

Generics are acquired early. Susan Gelman et al (2008) document the use of generics based on a longitudinal study on English-speaking children. The study was based on verbal interactions between the children and their parents. It involved eight children aged between two and three and followed them for over a year. Gelman et al found that generics were produced by even the youngest children. They also noted that children initiated generic talk and were not simply repeating their parents' utterances of generics. The rate at which generics were produced increased rapidly between the ages of two and four, at which point the children and their parents produced generics at the same rates.[106]

---

[103] See Leslie (2007:381).

[104] I'm speaking somewhat loosely about different types of generalisations as either sentences or their contents for better readability. However, to be clear, I'm assuming that sentences can express generalisations, which, as I argue in this chapter can also be the object of propositional attitudes. These sentences can either be overtly quantified, or generic and accordingly, express different kinds of generalisations: non-specific, generic ones or various other, more specific kinds.

[105] Example first given by Frederick (2005) and discussed in Kahneman (2011).

[106] Leslie takes this to mean that children use generics at the same rate as adults in general (2012:33). However, given that the research only examined child-adult interactions and not child-child or adult-adult interactions, I don't think this conclusion is warranted. There might well be additional reasons for why adults may use more generics than when speaking to children when to other adults some of which I speculate about in §5 of Chapter 5.

Susan Gelman and Lakshmi Raman (2003) conducted work that appears to show that the interpretation of generics is similar in children and adults. For example, children who used generics understood that they were not about specific individuals or instances and used similar linguistic and pragmatic cues to determine whether a sentence is generic or not. Importantly, the children, including the two-year-olds, already understood that generics tolerate exceptions. For example, in one of their experiments, participants were shown a picture of a member of a kind, which doesn't instantiate a typical property, such as a three-legged dog. When given suitable pragmatic and linguistic clues, for example, when asked whether *they*, not *it*, had four legs, both adults and children agreed that dogs have four legs.

Michelle Hollander et al (2002) conducted research which compared the understanding of generics with that of specific statements and overtly quantified generalisations in three and four-year-old children and adults. They found that the responses in English-speaking children, both three and four-year-olds, were comparable with those of adults when it came to generic questions. However, when concerned with overtly quantified generalisations containing 'all' and 'some', only the four-year-old children were adult-like in their responses, whereas the three-year-olds weren't. Instead, the three-year-olds responded to the overtly quantified generalisations as if they were generics.

Another study of note was undertaken by Leslie and Gelman (2012). In this study, Leslie and Gelman asked participants, again comprising of three and four-year-olds as well as adults, to remember facts about animals they had previously been given. These facts were either given in generic or overtly quantified forms. They found that children and adults had differing responses to the two quantified statements given (*all Fs are G*, *all of these Fs are G*). When the facts were given in generic form, they were recalled accurately by all age groups. However, both adults and children sometimes recalled overtly quantified generalisations as generic ones. Leslie uses this work to argue that generics cognitively primitive and therefore acquired earlier than overtly quantified generalisations.

2.4 Inferences

Generics don't only play an important role in teaching children about the world because they are acquired early, but also because they provide a fertile basis for inferences. Leslie draws attention to a phenomenon that she terms the 'generic overgeneralisation effect' ('GOG effect' from now). She describes the effect as 'the tendency to overgeneralise the truth of a generic to the truth of the corresponding universal statement' (2011:17). The GOG effect is meant to be an instance of our using System 1 processes, in this case to generalise, when System 2 processes would be more appropriate. A large number of studies have been cited as evidence of the GOG effect by Leslie. [107] [108]

[107] For example, Khemlani et al (2009, 2012), Meyer, Gelman and Stilwell (2011), Leslie et al (2011) and Leslie and Gelman (2012) are cited as evidence of the GOG effect.

[108] I won't be giving a broad overview and focus on particularly telling examples. See Lazaridou-Chatzigoga et al (2015, 2017, 2019) for a comprehensive treatment.

Sarah-Jane Leslie, Khemlani, and Glucksberg (2011) provided participants with sentences that were either generic generalisations (*Fs are G*), universally quantified generalisations (*All Fs are G*), or straightforward existential generalisation (*Some Fs are G*). Leslie et al found that participants frequently judged false universally quantified generalisations as true. In the examples, the properties ascribed are had by some, but not all members of the kind described. The proportion of false statements that were judged as true varied, depending on the content of the statement. Leslie et al compare responses to quantified analogues of what in Leslie's taxonomy are different types of generics: [109]

(1) All triangles have three sides (quasi-definitional)

(2) All pit bulls maul children (striking property)

(3) All tigers have stripes (majority characteristic)

(4) All lions have manes (minority characteristic)

(5) All Canadians are right-handed (false)[110]

Leslie et al found that a high proportion of universally quantified generalisations which ascribed a characteristic property were judged as true despite being false. Sentences in which the property is had by a majority of members of the kind, like (3), were judged as true 78% of the time, whereas sentences in which the property is had by a minority of members of the kind, like (4) were judged as true 51% of the time.

Leslie et al explain this phenomenon as follows: participants sometimes judge false universally quantified generalisations as true because they are treating them as their (true) generic analogues. For example, while (4) is false, its generic analogue, *Lions have manes,* is true and hence the fact that a significant proportion of participants judge (4) to be true can be explained if they are treating (4) as if it was *Lions have manes.* Leslie et al argue that this is evidence for Leslie's view that generics articulate System 1 generalisations, as she takes it to be characteristic of System 1 processes that they are sometimes employed when System 2 processes would've been more appropriate, in particular when cognitive demands are high. Leslie et al support their argument by arguing against three alternative explanations of their observations.

First, they investigated the possibility that participants were making wrong judgments out of ignorance. Leslie et al tested whether participants knew the relevant facts and found that they did.[111] Second, they considered whether participants understood the statements differently and thus gave answers that were

---

[109] For reasons discussed in Chapters 2 and 3, I disagree with Leslie's taxonomy, but will use Leslie's labels without qualifying here for readability.

[110] To clarify, (1) is the only true sentence in this list. The 'false' with respect to (5) is about the generic analogue of the sentence, *Canadians are right-handed.* According to Leslie, the generic analogue of (5) is false, whereas the generic analogues of (1)-(4) are true.

[111] Leslie et al tested whether participants judged generic ascriptions of the property to those members of the kind that don't have the property. For example, participants were asked whether statements like *Female lions have manes* was true. Participants judged these statements as false 84% of the time, which Leslie et al take to be evidence against the claim that ignorance plays a significant role in participants judging false universally quantified generalisations as true.

actually appropriate. For instance, they considered whether participants interpreted a statement like *All ducks lay eggs* to mean *All kinds/types of ducks lay eggs*. Leslie et al ruled out such interpretations by asking participants to paraphrase the generalisations in question. Barely any participants paraphrased the statements so as to suggest that they understood them as taxonomic statements. Third, Leslie et al considered whether participants understood the statements to range over restricted domains, for example whether they understood *All ducks lay eggs* as *All female ducks lay eggs* or even as *All female, fertile ducks lay eggs*.

Leslie et al ruled out this possibility by priming participants before they were asked to judge whether the statements in question were true. They did this by giving additional information that referred to the overall number of members of the kind. This was meant to make salient that the statement was to be read as being about all members of the kind and not some proper subset. Though the percentage of true judgements decreased, they remained at 60% and 30% for majority and minority judgements, respectively.

So, even after they are primed, a substantial proportion of participants make what at least appears to be a mistake, even if the majority don't. Having ruled out other explanations, Leslie et al take their observations to demonstrate that participants make generic generalisations by default which have to be consciously overridden to be corrected and hence supporting Leslie's approach.

3 Issues with Leslie's account

Leslie's account is closely tied to the dual-process approach. She posits generics as linguistic items specifically used for expressing cognitively basic generalisations, which are part of System 1. The data around generics, their acquisition, and prevalence in different types of speech and inference patterns are then expected to manifest characteristics in keeping with System 1 processes in the ways I discussed in §2 of this chapter. This has its benefits, but also some disadvantages.

Its benefits are that it neatly accommodates a theory of generics in a wider, popular framework of cognition with a wealth of empirical work that can be used to test its predictions, as Leslie does. In particular, her work on social generics can draw on Kahneman's idea that cognitive biases can arise as a result of inappropriate uses of System 1 processes.[112] However, this also means that Leslie's account is hostage to the fortunes of the dual systems approach. There may be good, more general reasons to be hesitant about the dual systems approach, but I won't address them here, and will focus specifically on the way Leslie's account relates to it. [113]

Because Leslie posits generics as expressing judgements made by a distinct Systems 1 faculty, she also hypothesises a clear dividing line between generics and overtly quantified generalisations both linguistically as well as with regards to the cognitive role they play. Although they are often portrayed that

[112] Leslie (2014, 2017), Wodak, Leslie and Rhodes (2015)

[113] A number of criticisms of dual process theory have been made. See Osman (2004) for in overview and one particular exchange chronicled by McCook (2017).

way, Systems 1 and 2 are not separate in either location or the processes they involve. Indeed, Kahneman (2002) himself points out that almost all processes are a mix of both (usually complementary) systems. This is reflected in the evidence concerning their acquisition and inference pattern. Contrary to Leslie's presentation of that data, it does not paint a conclusive picture of generic and non-generic generalisations as being made by processes that are neatly delineated and different in kind.

3.1 Acquisition again

As Lazaridou et al (2015) note, the broader empirical data on acquisition does not show a unified picture that would support Leslie's thesis. First, the studies mentioned above are the only two studies that found the differences between overtly quantified generalisations, specific statements, and generics. Other similarly designed studies do not show this effect. For example, Gelman and Raman (2003) record that the interpretation of generics as well as overtly quantified generalisations and specific generalisations is very similar in adults and young children, starting at 2 years old.

There are further reasons to be suspicious of Leslie's wide-ranging claims as a result of the overtly quantified generalisations that generics were contrasted with. As Lazaridou et al point out, just because children may have issues with some overtly quantified generalisations, this does not tell us much about whether they would have issues with others. Most of the overtly quantified statements used in the experiments cited by Leslie involve 'all' and 'some' but thinking that generics are quantificational does not mean that those would be the quantifiers involved.

James Kirkpatrick (2019) argues that the quantifiers that matter, A-quantifiers, are acquired at a similar age as generics. Kirkpatrick points out that the studies used by Leslie compare generics to overtly quantified generalisations containing D-quantifiers. Hence, even if the claim that D-quantifiers were acquired later than generics were substantiated, this need not tell us anything about when quantifiers are acquired and thus does not speak against A-quantificational approaches.

Kirkpatrick argues that A-quantifiers are acquired at a comparable time to generics. He cites various studies tracking when children acquire A-quantifiers. In one study by Antinucci and Miller (1976), Italian-speaking children are recorded as using A-quantifier expressions before the age of two. Another study by Cromer (1968) records A-quantifier expression usage in two-year old Polish-speakers. The data Kirkpatrick uses is limited in several ways. It is old, relies on small case numbers, and does not test quantificational adverbs that are particularly similar to what Gen is presumed to be like. So, though I agree with Kirkpatrick's point that the data used by Leslie, even if it was substantiated, does not tell against the idea that generics are quantificational, I don't think the data he uses is substantial enough to make the argument that generics are in fact acquired at the same time as quantificational adverbs in general. Nevertheless, his argument gives us reason at least to think that the available data does not support Leslie's view in particular.

3.2 Inference patterns again

Dimitra Lazaridou-Chatzigoga et al (2017) criticise Leslie et al's treatment of the alternative explanations. They argue that all three alternate explanations: ignorance, implicit domain restriction, taxonomic interpretations, as well as the atypical behaviour of 'all' can jointly explain the behaviour Leslie et al describe. I will focus on the latter two.

Lazaridou-Chatzigoga et al point out that 'all' might function in atypical ways that inform participants' responses to it in a manner that has nothing to do with generics. They argue that 'all' is potentially ambiguous between 'every' and 'very many'. This means that overtly quantified statements involving 'all' might be read as *Very many Fs are G* or *Almost all Fs are G,* and thus be judged as true simply because they are in fact true, at least in the case of some generalisations. This is an additional possibility that might explain some of the data, which is not explored by Leslie et al.

With regards to contextual domain restriction, Lazaridou-Chatzigoga et al argue that the primer might very well not be successful at countering contextual domain restriction. Instead, a primer that mentions exceptions to the generalisation, e.g., male ducks, would be more useful in checking whether participants understood the generalisations as contextually restricted.

Consequently, Dimitra Lazaridou-Chatzigoga et al (2019) conducted their own experiments in which they prime participants with explicit references to exceptions. To examine the plausibility of the contextual domain restriction, they used primers with different contextual information. These included primers that explicitly mentioned exceptions. Lazaridou-Chatzigoga et al found that with these contexts, acceptance rates decreased markedly for generalisations including 'all', 'all the', and 'each'. They argue that this suggests that domain restriction in combination with the other explanation can explain the evidence that Leslie takes to be in support of the GOG effect. Though further evidence is needed to substantiate their prediction, I do agree with Lazaridou-Chatzigoga et al that the data Leslie et al rely on does not provide strong evidence of a GOG effect, which in turn undermines one of the main points of evidence Leslie uses for her account of generics as giving voice to System 1 generalisations.

3.3 Cognitive demands

The hypothesis Leslie's account posed was not just that we make mistakes in which we use System 1 instead of System 2 processes, but specifically that we make such errors when the cognitive demands on us are great. However, as Lazaridou-Chatzigoga et al point out, the empirical work used doesn't tend to vary the cognitive demand participants are under.[114] Participants weren't required to respond in any particular amount of time or have any other tasks to do. Leslie et al simply assume that processing System 2 generalisations is cognitively more demanding than processing System 1 generalisations and thus that

[114] Lazaridou-Chatzigoga et al (2017:20).

participants sometimes default to System 1 generalisations, however, these assumptions are not independently motivated and, instead, part of her account of generics.

Meyer, Gelman, and Stilwell (2011) recorded the time it takes to judge the truth values of majority characteristic generics and overtly quantified generalisations featuring 'all'. They found that under time pressure, participants were more accurate (90%) at judging the truth of generics than they were at judging the truth of overtly quantified generalisations (55%). Participants were also faster at responding to generics (-900ms) than they were at responding to 'all' generalisations (-1000ms). At first sight, these results seem to support Leslie's classification for generics as dedicated expressions of fast System 1 generalisations.

However, Lazaridou-Chatzigoga and Stockall (2013) tested the time it took participants to read generalisations and the time it took to judge their truth-values for generics as well as generalisations involving 'all' and 'every'. They found a mixed picture. Some generics were much faster to read than the overtly quantified generalisations, but also than other generics. Specifically, they found that generics that would be classified as majority characteristic generics within Leslie's framework were much faster to process in this way. They found that when it comes to judging the truth-values for the statements, it took participants much longer to judge the truth-values of some generics, namely minority characteristic generics, than either other generics or the overtly quantified generalisations. Lazaridou-Chatzigoga and Stockall (2013) suggest that this might be due to participants having to engage in a process of quantifier domain restriction which would count against Leslie's view.

I find Lazaridou-Chatzigoga et al's points broadly convincing and therefore believe that the existence of the GOG effect is not well supported by current evidence. In combination with the mixed evidence in acquisition data, the empirical data for Leslie's proposal is inconclusive at best. Nevertheless, other evidence does suggest that generics are inferentially fertile in other ways that should be explained by an account of genericity in mental representations.

Cimpian, Brandone et al (2010) report an intriguing discrepancy between the prevalence a property must have among members of a kind to be accepted, and the prevalence estimated by people who believe the generic in question. They conducted a study in which participants were assigned one of two groups. The first group was presented with a list of generics and facts about which proportion of members of the kind have the property.[115] For example, participants were told that 30% of lorches, a made-up animal kind, had purple feathers. They were then asked whether the corresponding generic, *Lorches have purple feathers* was true. The second group was presented with the same list of generics and then asked to estimate what proportion of members of the kind, e.g., lorches, have the relevant property. The same experiment was repeated with a 'most'-quantified generalisation instead of a generic.

[115] Cimpian et al (2010:1457)

Cimpian et al report that when participants were given a 'most'-quantified generalisation, the average level of prevalence at which participants in group 1 accepted the 'most'-generalisation was 76.9%, very similar to the average estimates given by group 2 at 78.0%. In contrast, the average level of prevalence at which the generic was accepted was lower, at 69.1%, but the average prevalence estimated was much higher, at a staggering 95.8%. So, while generics can be accepted at relatively low prevalences, forming a generic belief leads us to estimate very high prevalences.

These experiments were first conducted with generics concerning novel kinds so as not to have people's background beliefs distort their judgement.[116] However, Khemlani, Leslie, and Glucksberg (2012) repeated them using generics concerning kinds that participants were familiar with and found similar results. This indicates that there are distinctive features of generic mental representations, specifically with respect to inference patterns they give rise to, that need to be accounted for by theories of genericity in mental representations.

Much of the evidence used by Leslie to support her view is inconclusive. At this point, I think that the empirical evidence simply isn't sufficient to motivate her approach, but neither does it conclusively show that Leslie's account is ill-fated. I will therefore move on to a theoretical issue with Leslie's account of genericity in mental states that I want to explore briefly before proposing an alternative view.

3.4 Theoretical issues

I briefly described Leslie's theory of the facts that can make true different kinds of generics in Chapter 2. Leslie holds that, as generics give voice to System 1 generalisations, the types of facts that can make them true correspond to the conditions under which our System 1 generalisations are accurate. For example, according to Leslie, generics that ascribe dangerous properties can be made true by weak dispositional links even if very few members of a kind actually have the relevant property. This is *because* our System 1 processes are quick to make generalisations when the property ascribed is dangerous.

However, Leslie does draw a distinction between the types of facts that can make a generic true and the conditions under which it is true. She calls the former the 'metaphysical' and the latter the 'semantic' truth-conditions of generics.[117] The semantic truth-conditions are what we specify as part of a semantic theory and also what is involved in linguistic competence. The metaphysical truth conditions are descriptions of the types of facts that make true the generics and encode the psychological salience of properties.

Leslie simply disquotes the semantic truth conditions for generics.

---

[116] Although the participants did not have prior beliefs about lorches themselves, they likely did have other background beliefs, for instances about the distribution of a property like having feathers amongst the kinds of things that have feathers, i.e. birds. I discuss how such background beliefs influence the inferential patterns generic beliefs give rise to in §5 of this chapter.
[117] See Leslie (2008:43)

(6) Tigers have stripes

According to Leslie, (6) is true iff tigers have stripes. She does not provide a further, informative account of these conditions. However, she does give an informative account of the truth-makers for (6), namely the fact that most tigers are striped. For Leslie, the metaphysical and semantic truth conditions need to be distinct as the metaphysical truth-conditions for generics are too complex to be expressions of cognitively basic generalisations and thus couldn't explain how they are acquired at an early age. However, I think that this distinction means Leslie deprives herself of an important theoretical tool: to explain the (semantic) truth conditions of generic contents in terms of their cognitive role or vice versa.[118]

Even though for Leslie, generics articulate default generalisations, the truth conditions that we as language speakers would learn are not the conditions under which our generalising faculty is accurate. This disconnect between the conditions under which generics are true and the facts that make them true is problematic. First, it seems very strange to me that the conditions under which a sentence is true would not be identical to (at least some of) the facts that make them true. Truth conditions for a sentence should describe whatever it is that makes them true. Second, it deprives Leslie of the theoretical tools to explain how we can use generics with the ease and proficiency from an early age that she emphasises in virtue of what generics mean.

This disconnect is unnecessary and disposing of it, I will argue, opens the door to understanding the linguistic role of genericity in terms of the mental states generics express and vice versa, something I will pursue in the next chapter. In the second part of this paper, I propose a view of generic beliefs that identifies the semantic contents of generic sentences with the contents of generic beliefs and thus does not posit a disconnect between genericity in language and thought.

4 Generic belief

4.1 The invariant version

The alternative approach consists in extending the generality account I proposed in Chapter 3 to an account of generic beliefs. The core idea of the generality approach was that generics semantically encode non-specificity by existentially quantifying over non-generic generalisations. In the previous chapter, I defended the invariant version of the account against over-generation worries. However, in this and future chapters, I will draw out some of the ways in which the contextualist version of the generality provides an impoverished picture of the role of generic belief. I start by explaining the invariant account of generic beliefs before turning to the contextualist version.

According to the generality account, a generic of the form *Fs are G* is true iff at least one of several non-generic generalisations about Fs and being G is true. Because the generality account in its invariant form

[118] For other, to my mind valid, criticism of Leslie's distinction see Assarsson (2012) and Sterken (2015b).

explains the meaning of generic sentences in terms of the non-specific generalisations they express, it is straightforward to extend it to an account of generic beliefs. The same propositions that can be expressed by generic sentences can also be the objects of propositional attitudes, like belief. So, for instance, a belief that Fs are G is a belief in a proposition of the same type that could be expressed by the generic *Fs are G*. I will use an example to illustrate this picture of generic beliefs.

(7) Wolverines are aggressive

I recently watched a survival TV programme in which one of the contestants encountered a wolverine at night and killed it. Before this, the only belief I had about wolverines was that they eat a lot, due to a quirk of false etymology.[119] To justify the contestant's actions, the programme narrator said: 'Wolverines are aggressive'. The programme didn't provide any information about whether wolverines displayed aggressive behaviour particularly often, are only disposed to it under particular circumstances, or were simply more aggressive than similar animals, like other weasels. On the basis of the information provided, I formed the belief that wolverines are aggressive. The belief that wolverines are aggressive was non-specific, it wasn't an attitude towards a proposition of any specific quantificational force or flavour. The generality account of generic beliefs can explain how this is the case as the proposition I took an attitude of belief towards is itself non-specific. However, the account may fare less well when considering other generic beliefs that seem like they might be more specific.

(8) Ticks carry Lyme disease

Contrast my belief that wolverines are aggressive with my belief that ticks carry Lyme disease. I know of several non-generic generalisations about ticks and Lyme disease that are false. For example, I know that most ticks don't carry Lyme disease, that several subtypes are incapable of carrying it, and that there is no social norm according to which they should carry Lyme disease. Instead, I know that some ticks do carry Lyme disease and that some types are disposed to carry it under certain circumstances, though I don't know much about what these circumstances are. The situation then is that I do believe that at least one non-generic generalisation about being a tick and carrying Lyme disease is true. However, given that I have additional beliefs about the nature of this generalisation, my belief that ticks carry Lyme disease seems much more specific than my belief that wolverines are aggressive.

According to the invariant view, my belief that ticks carry Lyme disease must still be similarly non-specific: it is still just a belief that at least one non-generic belief about ticks and having Lyme disease is true. The proponent of the invariant generality account can explain the contrast to my belief about wolverines by arguing that in addition to the non-specific beliefs about ticks and Lyme disease, I also have *additional*, more specific beliefs. So, for example, I believe both that ticks carry Lyme disease, and also that

[119] Names for wolverines in many European languages ascribe gluttony. For example, wolverines are called 'glouton' in French, 'Vielfraß' [eats-a-lot] in German and rasva-maha [fat-belly] in Finnish. This is most likely based on an initial mistranslation of 'fjellfross', meaning mountain cat, in Norwegian. See [https://www.duden.de/rechtschreibung/Vielfrasz]

most ticks don't carry Lyme disease. These in turn give me other beliefs about which generalisations, which could make true the generic, are not the ones that do make it true, for example the generalisation that most ticks carry Lyme disease.

This is a familiar phenomenon. My friend Paul is currently doing his PhD in the United States. Based on me telling you this, you will have, given that I am reasonably reliable and so on, formed the belief that Paul lives in the US. However, if I now tell you that Paul lives in New York City, you will form an additional belief, that Paul lives in New York City, which builds on, but is separate to, your prior belief that Paul lives in the US. That prior belief hasn't gone, as you still believe that Paul lives in the US. However, you now also have additional beliefs about where Paul lives. Another reason to think that these are different beliefs is by noticing that the dispositions that they give rise to, or consist in, are different. Before I told you where in the US Paul lives, you wouldn't have known where to catch a plane to visit him, whereas now you do. However, some might still want to maintain that my belief that ticks carry Lyme disease is more specific than my belief that wolverines are aggressive. One option for them is to become a contextualist.

4.2 The contextualist version

Here is an alternative picture of generic beliefs which can say that some generic beliefs are more specific than others, based on the contextualist version of the generality account that I discussed in the previous chapter. On the contextualist version of the generality account, generic sentence types are non-specific, but generic tokens are of varying degrees of specificity. In particular, the contexts in which generics are uttered determine how many and which non-generic generalisations are salient. Contextualism of any kind is a theory about linguistic items, not the beliefs they can be used to express. So, acknowledging that contextualists are primarily interested in language, what *could* they say about the beliefs generics are used to express? I will discuss the contextualist version of the generality account only, not other versions of contextualism here, although I will be calling the proponent of the contextualist generality account 'contextualist' for brevity.[120]

For contextualists, what I hold when I believe that I am hungry or that East Anglia is flat are beliefs that involve components that can be the referents of context-sensitive expressions like 'I' and 'flat'.[121] Nevertheless, the beliefs themselves clearly are about me, or ascribe a flatness with respect to a standard. Similarly, a contextualist about generics could think that generic beliefs are beliefs that at least one of

[120] Other contextualists will have to give an account of the beliefs non-specific uses of generics express and can give rise to in others. This would, along the lines of argument considered in Chapter 2, take the form of either thinking that such beliefs are attitudes towards no, or several beliefs. I consider why this might be difficult in Chapters 5 and 6.

[121] I am assuming contextualism about 'flat' and 'I' here, but my point isn't specific to this expression and could be made in the same way about any other context-sensitive expression.

several contextually salient generalisations about a kind and a property is true. In this sense, generic beliefs token generic types, just as utterances can.

In the wolverine example, the contextualist could say that I formed a belief that one of the contextually restricted set of generalisations about wolverines and being aggressive is true. However, the set of eligible generalisations is relatively unrestricted, and hence so is my belief. In contrast, because I already have background assumptions about ticks and Lyme disease, I might still form a belief that ticks carry Lyme disease that I would express using a generic, but the belief is much more fine-grained that my wolverine belief.

The above contextualist sketch might raise some questions about the individuation of generic beliefs.

(a) Can generic beliefs be as specific as belief in non-generic generalisations?

(b) Can one and the same generic belief become more or less specific?

(c) Can two people who believe that dogs bark have different beliefs?

The contextualist answer to (a) is yes. Beliefs in highly specific generic propositions can be identical to beliefs in non-generalisations. For instance, a belief that most dogs bark can be expressed using either a generic or the overtly quantified 'Most dogs bark'. Any belief that could be expressed using a generic is thereby a generic belief, even if it could also be expressed using an overtly quantified generalisation. The purpose of calling such beliefs generic beliefs lies in emphasising their structural similarities to less specific generic beliefs. The important thing about generic beliefs for the contextualist is that they can be non-specific, not that they always must be.

The answer to (b) is no. A belief that dogs bark really is a belief that of a fixed set of non-generic generalisations at least one is true. If the set of non-generic generalisations is smaller in one belief than for another, but their content is otherwise the same, say if they are both about dogs and barking, then the beliefs are similar but not identical because the relevant beliefs are attitudes to different propositions.

For the same reasons, the answer to (c) is yes. Consider other beliefs that that can be expressed using context-sensitive language. My belief that my brother isn't tall for a basketball player is not in conflict with my belief that he is tall for a European man, because they are two distinct beliefs. That's the case even though they can be expressed using sentences that look like the negation of one another. Likewise, two people can believe that my brother is tall without believing the same thing, e.g., if one believes that my brother is tall for a basketball player and the other that he is tall for a European man.

4.3 The basic view again

Though I have described the contextualist and invariant versions of the generality account as two alternative options, I believe that the predictions they make about individual cases are similar. Whereas the invariant account predicts that I believe in several specific generalisations about ticks and Lyme disease *in addition* to a non-specific generalisation, the contextualist account only posits one belief in a

generalisation of a particular level of specificity. However, thinking of beliefs not, as Hawthorne puts it, a 'box that contains sentences of English' (2004:109fn), but, instead, of as dispositions, the dispositions each theory predicts look equivalent. Where the accounts differ is in whether they afford a distinctive role for generic beliefs.

Compared to the invariant view, the contextualist account paints a fragmented picture of generic beliefs. If generics express the same kinds of generalisations that could also be expressed using overtly quantified generalisations, then there are no distinctive generic beliefs. The beliefs that can be expressed using generics are themselves generic only insofar as they can be expressed using generics. By contrast, the invariant version of the view, posits distinctive non-specific generalisations as the objects of distinctive belief. I take this to be an advantage for the invariant version of the generality account. I will argue in Chapter 5 that generic generalisations play a particular epistemic role as the objects of evidentially undemanding beliefs and consider a subset of such beliefs, ones about social kinds, in Chapter 6.

The picture I have sketched in this section is a basic extension of the generality account. It simply says that the same generic contents that are specific to varying degrees can be believed as well as expressed. I have said little about what it takes to hold a generic belief or take any other propositional attitude towards it. In the next two sections, I will discuss the inference patterns generic beliefs feature in. I will explore this aspect of generic beliefs, what they dispose us to, in the next section.

5 Generic inferences

Beliefs are often analysed in terms of dispositions. My belief that the Chinese supermarket below my flat opens on Tuesdays either is or is at least evidenced by a bundle of dispositions. For example, in executing my plans to go to the Chinese supermarket, I depend on my belief and hence go on Tuesday instead of Monday. I would also be surprised should it turn out that the Chinese supermarket is actually closed on a Tuesday. Thus, to better understand generic beliefs it will be useful to look at what they dispose us to.

Part of the way in which generic beliefs play an important role in cognition is through giving rise to other beliefs in inferences. Given what I've said so far, you might expect that the generality account predicts that generics don't have particularly strong inferential powers precisely because they are so non-specific. There's not much someone can draw cognitively from knowing that they feel a bit ill, definitely not as much as if they knew exactly what was wrong with them. In fact, it seems like the flip side of non-specificity is that there is not much information to draw on to inform one's actions. However, this seems to be in conflict with the empirical work I drew attention to earlier, which suggests that generic dispose people to draw very strong inferences. In this section, I want to diffuse this tension by making use of the notion of defeasible inferences to explain how generic beliefs can be inferentially powerful despite their non-specificity.

5.1 Defeasible inferences

Francis Pelletier and Nicholas Asher (1997) focus on inductive inference patterns in which generics feature. They argue that generics license (in a normative and descriptive sense) a variety of defeasible inference patterns, including one they call 'Commonsense Entailment'.[122] Commonsense Entailment ratifies the defeasible modus ponens, an inference of this form:

(9) (a) Dogs bark

(b) Vera is a dog

(c) Vera barks

This inference is defeasible because for instance, on learning that Vera is mute, one is no longer licensed to infer (9c). This inference pattern seems well-suited to explaining the inferential power of generic beliefs that the experimental studies mentioned in §2.4 documented.

I want to explain why this type of inference pattern exists in a way that is compatible with the generality account of generic beliefs. The rough idea that I already mentioned in §7 of Chapter 2 is this: even though the belief that dogs bark deductively licenses us only to infer that at least one non-generic generalisation about dogs and barking is true, our cognitive biases and long-standing background beliefs lead us to expect several non-generic generalisations to be true or false at the same time, which in turn leads us to believe more specific generalisations and beliefs about individual members of the kind.

I want to start with a somewhat speculative story about why this might be. Though different non-generic generalisations have different truth-conditions, they are often true or false at the same time. It is true that most dogs bark, but also that they are capable of barking, disposed to it, and that barking is characteristic of them.

Not only are they often true and false at the same time, but this is also the case because they often explain one another. Most dogs bark because dogs are capable of and disposed to barking. Even if the generalisations don't always explain one another, we often think that they do. For example, we often believe that things have essences and that these ground both normative and other properties independently of whether this actually is the case. On the basis that, say, most dogs bark we might infer that this is so as a result of their shared *nature*, and even that this in turn makes it such that dogs *should* bark and hence that mute dogs are in some sense defective.

These explanatory patterns can explain why speakers have general expectations of co-occurrence about different non-generic generalisations. In believing that dogs bark, we believe that at least one generalisation about dogs and barking is true. However, we also have the expectation that if one generalisation about dogs and barking is true, others will be too. If on the basis of this expectation we

[122] Pelletier and Asher (1997:32)

then believe that some prevalence generalisation is true, we might well also expect a given member of the kind to have the property in question.

These co-occurrence expectations patterns are not general. On learning that many teachers suffer from stress, we don't therefore infer that all of them do. Co-occurrence expectations are not general and can be constrained by prior beliefs. For example, recall my discussion of *ducks lay eggs* in the previous chapter. Because we believe that most ducks don't lay eggs, on account of less than half being fertile females, our belief that one generalisation about ducks and the laying of eggs is true doesn't generate as many co-occurrence expectations as generic beliefs about kinds and properties about which we have fewer prior beliefs. Instead, which co-occurrence expectations we have seems to me to be likely to be influenced both by general cognitive biases and specific background beliefs concerning both the kind and property ascribed.

5.2 Cognitive biases

As for the cognitive biases which explain these expectations of co-occurrence, we can make use of Gelman (2003) and Leslie's (2007) work on essentialising. Gelman posits a general cognitive essentialising bias, which causes us to assume that superficial properties had by members of a group are had in virtue of an underlying, stable essence. Leslie argues that generics activate this bias.[123] According to Leslie, when we learn information via generics, we tend to assume that superficial properties ascribed in the generic are had on the basis of some deep properties members share, as a result of this essentialist bias. Once we have essentialised members of the kind in this manner, we also expect them to be somewhat uniform with regards to their superficial properties, such that if some members of a group have a property, others are likely to have it too.

The work Gelman, Leslie and others have done is used in Leslie's account, but it can also be combined with other accounts. In particular, the generality account can posit essentialist or other cognitive biases to explain expectations of co-occurrence. The essentialist bias is only one of several strands of explanation required here and a more substantial explanation would be needed for a fuller account of which defeasible inferences generic beliefs dispose us to.[124] However, the explanation offered in this subsection shows that the generality account has the potential to explain why generic beliefs, though they themselves are non-specific, dispose those who hold them to co-occurrence expectations towards the different non-generic generalisations that could make the generic belief true. As a result of the co-occurrence expectations they,

[123] Hoicka et al (2018) conducted a study that suggests that generics are not unique in causing essentialisation. This is compatible with my explanation here.

[124] For instance, I described the essentialist and normative biases Haslanger (2014) posits in the previous chapter. I also argued that relevance expectations can constrain which beliefs, in addition to non-specific generic beliefs, we form on hearing a generic asserted.

in the absence of prior conflicting beliefs, form beliefs in non-generic generalisations and expect given members of the kind to exhibit the property in question.

6 Conclusion

In this chapter, I proposed a new account of generic belief. The proposal connects semantic accounts of generics to genericity as a feature of mental representations. It does this in a straightforward way: by arguing that the propositions expressed by generics are also the objects of generic beliefs.

I motivated this account by comparing it to Leslie's approach, according to which generics articulate default generalisations. Empirical support for Leslie's view is limited and it involves an awkward disconnect between the semantics of generics and their mental role. I argued that the account I propose is simpler and doesn't feature the same disconnect between mental and linguistic representation.

I developed a particular proposal by extending the generality account that I argued for in the previous chapter. On this account, generic beliefs are non-specific. They are attitudes towards the same non-specific generalisations that I argued in the previous chapter are the semantic values of generic sentences. I suggested that the generality account can give a plausible picture of both coarse and fine-grained generic beliefs even in its invariant form and the inference patterns they give rise to. In the next chapter, I will utilise the account of generic beliefs I have proposed to consider the functional role of generic generalisations in thought and language.

Chapter 5: What generics are for

1 Introduction

In previous chapters, I presented a novel account of genericity in language and thought. In doing so, I have mainly been focused on the content of generic sentences and beliefs. In this chapter, I want to turn to a different question: what is genericity for? In other words, what needs do generic tokens, in particular generic beliefs and sentences, help us satisfy? I will identify one function that I believe to be central to understanding the role of genericity in thought and language: that generic tokens enable us to believe, know, and communicate *evidentially undemanding* generalisations, by dint of their non-specificity.

I start the chapter by setting out what I mean by saying that generic beliefs and sentences have functions and why I take a characterisation of these to be worthwhile. I respond to the objection that linguistic constructions, including generics, do not have any proper functions beyond expressing contents in §2. I then turn to my core argument. In §3, I argue that the picture of generic beliefs I presented in Chapter 4 suggests that generic beliefs are evidentially undemanding: they can be supported by a large variety of evidence. In §4, I argue that this means that we often have good evidence for a generic, even when we don't have good evidence for any one non-generic generalisation. This makes generics useful: they enable us to form inferentially powerful beliefs on the basis of limited evidence. I then turn to the significance of such beliefs in our conceptual apparatus. Following Johnston and Leslie (2012), I argue in §5 that much of our knowledge about kinds is generic in form. Consequently, seeing how the non-specificity of generics makes them epistemically attainable for us can help our understanding of how we are able to categorise things in the world. In §6, I return to generics in their linguistic guise. I argue that generics enable us to express generic beliefs by virtue of the fact that those beliefs are easy to gain epistemic justification for. This makes them an interesting instance of hedging speech.

2 Functions

Most things have functions. Watering cans enable people to direct water on to things like soil. Our legs enable us to walk. Often, understanding what something is consists, at least in part, in understand its function(s). For instance, if we knew only about the physical properties of legs, but not that they are for walking, then we would have quite a limited understanding of what legs are. This isn't just the case for biological entities and artifacts. For example, money facilitates trading, wakes facilitate grief, and language facilitates the sharing of information.

2.1 Functionalist accounts

Many have been tempted to give philosophical analyses of what it is to be a certain type of thing by ascribing a function to it. For example, Edward Craig (1999) argues that a social function of attributing

knowledge is to mark out potential reliable witnesses by ascribing knowledge to them. Understanding this, according to Craig, is partly what it *is* to understand the concept of knowledge. Likewise, functionalists about mental states believe that what it is to be a certain type of mental state simply *is* to fulfil a particular role, i.e., to stand in causal relations to sensory stimulations or behaviour.[125] According to functionalists, what realises mental states doesn't tell us much if anything about what they are. What matters is what roles they fulfil.

We don't need to think that to be informative, functional accounts need to make claims about there only being one central function for each concept. For example, though Craig (1999) draws attention to one social function of knowledge, others have presented additional functions, which are meant to be compatible with Craig's account.[126] Things can have several functions, which can be independent or derivative of one another.[127]

I believe that generics are likely to have several functions some of which are contingent on features of the natural language in which they are expressed. For instance, in English, generics are usually shorter than sentences that express non-generic generalisations making them useful as means of communication when one has limited time or space to convey one's message, such as in newspaper headlines. However, in this chapter I want to focus on one particular function of generic beliefs that I take to be important to understanding genericity cross-linguistically and in both thought and language: evidential omnivorousness.

In addition to giving an insight into the nature of the things characterised, considering their functions can also help constrain other theorising about them. For example, if the function of a mental state is to play a specific causal role, then so much the worse for theories that are incompatible with the identified function. Functional talk can also help us better understand the phenomena in which the thing in question exhibits proper or improper functioning. For example, functional analyses of knowledge can help us understand how we signal that some people are reliable sources of information, but also part of what goes wrong when people are wrongly denied their status as knowers.

My hope in this chapter is that characterising a function of genericity does both. First, it can constrain our theorising about generic tokens in the sense that a theory of genericity should be able to account for the ways generic tokens occur. To the extent that non-specificity approaches to genericity, like the generality account, fulfil this requirement, I take that to speak in their favour. Second, I hope that characterising the epistemic and communicative roles genericity plays can help us better understand various cognitive and communicative phenomena. I will apply the argument from this chapter to one such phenomenon, stereotyping, in Chapter 6.

---

[125] For instance, see Putnam (1963), Lewis (1980), Shoemaker (1981).

[126] For example, Kusch (2009), Beebe (2012), and McGrath (2015), although others have also proposed alternative functions. See Kusch and McKenna (2020) for an overview.

[127] I discuss one such secondary function in §6 of this chapter.

So far, I have talked about different types of things having functions: artifacts, biological entities, but also concepts, like knowledge. In this chapter, I will be talking about genericity as having a function. Cappelen has recently criticised the use of function talk in our analysis of language. His reasons, if successful, would also spell trouble for functional talk about generics as linguistic items.

2.2 Cappelen against functions

Herman Cappelen (2018) criticises the use of function talk in the conceptual engineering literature. Though he is concerned with the function of concepts, or rather, individual words denoting them, his point applies to talking about functions of linguistic constructions more generally.

Cappelen doesn't think that words have any uniform function other than to pick out referents. Using the example of 'freedom', He points out that people use the term to do many things, such as promote freedom, undermine it, or celebrate it. The uses of the word 'freedom' and, in particular, the intentions of speakers who use the term, are too varied to have one informative and uniform function beyond referring to freedom. Cappelen says: 'The reason "freedom" is useful is that it can be used to talk about freedom' (2018:187).

I do agree with Cappelen's broad point, namely that it seems very hard to identify one intention, beyond the narrow communicative intention of referring, as the purpose of any one speech act. Although I did note earlier that things can have several functions, independent or derivative ones, it does seem reasonable to require that not just any purposes for which something can be *used for* is therefore its function. Instead, we should be looking for broader patterns of use, and I take it to be Cappelen's point that, very often, those terms that interest conceptual engineers don't feature in such broader patterns. However, I don't think Cappelen's point invalidates the use of functional talk about genericity, or indeed many other phenomena.

First, my focus will be on genericity as something that can come in linguistic, but, importantly, also cognitive guises. My argument in the first part of this chapter will be that non-specific generalisations do have distinctive uses as the objects of mental states that arise out of their non-specificity. This part of my argument is thus insulated from Cappelen's charge insofar as it applies only to linguistic items.[128]

Second, I don't think that the fact that people can have different intentions when using 'freedom' shows as much as Cappelen suggests. For one, even if we did accept that functions for linguistic items do have to include some intentional component, we can distinguish between stable functions of a concept and what's sometimes called accidental uses of it.[129] For example, a chair can be used to sit on, but also to sustain bonfires, allow people to reach high-up things on a shelf, or lift people in the air at celebrations.

---

[128] There may be a belief-based version of Cappelen's point, namely that the basic function of generic beliefs just is to enable us to form beliefs about whatever propositions generics express. I agree with this, but don't think it invalidates talking about functions for reasons I explain in the next paragraph.

[129] Preston (1998)

The fact that a chair can be used for many things, doesn't mean that it doesn't have a basic function: to seat people. More importantly, however, functional analyses about linguistic entities don't have to be about speaker intentions at all.

Thomasson (2020) suggests that can understand a linguistic item's function by asking 'what we can do better with such a vocabulary than if we lacked it'.[130] She says:

> If we think of a language, and the terms in it, as human creations—as artifacts […], it is natural to think that linguistic items, too, may have proper functions, identifiable separately from the diverse uses to which they are put. I have elsewhere […] distinguished the 'practical significance' or proper function of a range of vocabulary (the function it serves in our overall linguistic apparatus, which explains why it is useful to have vocabulary like that in place), from that of its use in different speech acts (what it is used to do on particular occasions). Even where uses vary, a more stable proper function or functions may be identified. (2020:446)

It's this practical function as Thomasson describes it that I am interested in. Generic generalisations, in thought and speech, enable us to satisfy certain needs that we would not be able to meet if we didn't have non-specific generalisations as the objects of belief or the semantic contents of sentences. It is true that generic generalisations *now* have this capacity in virtue of being objects of belief or semantic values of sentences. But this doesn't preclude talking about what needs they enable us to satisfy, and the way they do so, even the latter explanation is directly about the objects of generic beliefs, or the meanings of generic sentences.[131]

In what follows, I will focus on one function of genericity that I take to be central in this regard. The function in question is in enabling people to make rough generalisations, which are evidentially undemanding, but inferentially powerful in a manner that affords them a central role in cognition and, derivatively, in language. Hence, I start by looking at genericity in thought.

3 Generic beliefs as evidential omnivores

In the previous chapter, I suggested that a straightforward way of understanding generic belief is as an attitude towards those the non-specific generalisations expressed by generic sentences. Generics are linguistic items that express propositions, generic propositions. Such propositions can be expressed linguistically, but we can also take a variety of attitudes towards them. We can hope for them, dread them, expect them, or we can *believe* them. Such beliefs are generic beliefs.

In the rest of this chapter, I will argue that the non-specificity of generic beliefs has one particular epistemic function, which can in turn illuminate their cognitive and communicative role. The first step in

---

[130] This formulation does not appear in Thomasson (2020) but is quoted by Cappelen (2018) from an ancestor of Thomasson's paper.

[131] Thanks to Matthieu Queloz for discussion. See Queloz (2021) for discussion of functional analyses.

my argument will be to argue that generic beliefs are evidentially undemanding. The idea that generic beliefs are evidentially undemanding is a consequence of the idea that they are non-specific.

3.1 Evidential omnivorousness

If generics express non-specific generalisations, then generic beliefs inherit this non-specificity from the propositions they are attitudes towards. This means that generic beliefs are easily made true, just as other non-specific beliefs are. The belief that Miriam went to a supermarket is more easily made true than the belief that Miriam went to Morrisons because all the facts that would make the latter true would also make the former true, but given that there are non-Morrisons supermarkets, the converse doesn't hold. Likewise, a belief that bats have good hearing can be made true by the facts that would make true the belief that most bats have good hearing alongside facts that make true a variety of other non-generic generalisations.

The fact that generic beliefs are easily made true also makes them evidentially undemanding. Because generics are non-specific, evidence for any non-generic generalisation about a target kind and property is also evidence for the relevant generic. Just as evidence of Miriam having gone to Morrisons is evidence for my belief that she went to a supermarket, evidence for a more specific, non-generic generalisation is also evidence for the corresponding generic belief.

Say I have evidence for believing that most bus drivers are grumpy. Given that my belief is made true by any of several specific generalisations, the evidence of my belief that most bus drivers are grumpy is also evidence for my belief that bus drivers are grumpy. Likewise, if I have evidence that bus drivers share some common characteristics that disposes them to be grumpy or that they are grumpier than some relevant comparison class, say truck drivers, then this evidence also acts as evidence for the belief that bus drivers are grumpy. Generic beliefs are evidential omnivores: they take whatever evidence they can get.

3.2 Easy justification

Their evidential omnivorousness makes generic beliefs highly useful. This is because they maximise the number of justified beliefs we can form on basis of evidence available to us.

First, being evidentially omnivorous means that there is a large variety of types of possible evidence that can justify a generic belief. As long as we have evidence for believing any of several non-generic generalisations, we have evidence for the relevant generic. Some facts are statistical, some dispositional, and some normative. In a basic sense, if evidence for any of a set of non-generic generalisations is also evidence for the corresponding generic, but not the other way around, then the generic belief is more easily justified than the non-generic generalisations.

However, it's not just that we have evidence for a generic belief whenever we have evidence for a corresponding non-generic generalisation. Generic contents also enable us to form justified beliefs when we have evidence that wouldn't support *any one* non-generic generalisation. We can have evidence for

thinking that at least one non-generic generalisation is true without that evidence being evidence for any one of them individually. This type of knowledge is a hallmark of knowledge with non-specific content. For example, I can know that a dog chewed up my ball, say by looking at the bite marks, without knowing what kind of dog it was. I can know that I have a missed call from a relative by recognising the area code of the city they live in, without knowing which one it was that called me. The same applies in the case of generic beliefs.

Let's go back to the wolverine example from the last chapter. Remember that I formed the belief that wolverines are aggressive based on watching a TV programme. Here are some non-generic generalisations that I could have believed instead:

(1) Most wolverines are aggressive

(2) All wolverines are disposed to act aggressively under certain circumstances, e.g., if encountering possible predators when with their young

(3) Wolverines display higher levels of aggression than some comparison class, e.g., martens

Given that I was fairly ignorant about wolverines, I didn't have much evidence for any of (1)-(3). I didn't know how widespread aggressive behaviour is among wolverines. I didn't know whether it is a general behaviour or only occurs under specific and limited circumstances. And I didn't know anything about how the rate of aggressive behaviour in wolverines compares to that of other weasels. It seems to me therefore that I didn't have any good evidence for believing any of (1)-(3). However, I did have good evidence for thinking that at least one of (1)-(3), or some other similar non-generic generalisation, is true. In the next section I detail what forms of evidence we can have for a generic belief in the absence of evidence for corresponding non-generic generalisations.

4 Making the most of what evidence we have

4.1 Non-specific evidence

In this section, I argue that generic contents, by dint of their evidential profile, enable us to form justified beliefs based on limited evidence and that this is a central function of theirs. Before I describe several such types of evidence, a note of clarification. Although I will be speaking of generic generalisations as enabling us to do things, I in no way mean to imply that we consciously choose to form them. We don't have that kind of control over what to believe. Instead, my argument is that generic generalisations, as objects of generic beliefs, fulfil a certain role that is important to understanding them.

4.1.1 Testimony

Testimony is an important source of justification. Often, being told that things are a certain way gives us justification for believing that they are that way, or even know that they are that way.[132] However, information conveyed via testimony is often itself unspecific. Sometimes, our informants themselves use generics to convey a generalisation, as was the case in the programme on which I based my belief about wolverines. In this case, the TV presenter simply used a generic, *Wolverines are aggressive*. Given that I received no information beyond the generic, I had no good evidence for believing any more fine-grained generalisation. But I did have good evidence for believing that wolverines are aggressive.

4.1.2 Memory

Another type of justification for non-specific beliefs occurs when we fail to remember the details of a piece of information. For example, we might have had access to detailed information in the past, whether via testimony, perception, or other means, but only remember broad-brush elements.

Imagine you have in the past learnt that slipped disks are comparatively prevalent in bus drivers. However, while you remember with close to certainty that there is a link between being a bus driver and having slipped disks, you've forgotten whether the link was that (i) most bus drivers, sooner or later, suffer from slipped disks; (ii) that bus drivers are disposed to slipped disks, by virtue of the work they do; or (iii) whether the link was comparative, such that more bus drivers than truck drivers suffer from slipped disks. In this case, you wouldn't have any good evidence for believing any of the more specific generalisations, but you would have good reason for believing the non-specific generalisation.[133]

4.1.3 Symptomatic evidence

The last kind of evidence, and one that I think is especially important, is what I want to call 'symptomatic evidence'. This kind of evidence is information that would, in combination with other information, provide evidence for a specific generalisation, but where we do not have access to that other information. For instance, we often observe instances of members of a kind who have some property. Even when we have observed several such individuals and no individuals that lack the property, we often fail to have the additional evidence that would justify forming more specific generalisations because we don't know how representative the individuals we encounter are, and only observe them in a limited set of circumstances.

For example, say I encounter several wolverines that behave aggressively. Now, without knowing whether the wolverines I have encountered are representative of wolverines more generally, I don't have particularly good evidence for believing that most wolverines are aggressive. Likewise, unless I know how prevalent aggressive behaviour among some salient comparison class is, I don't have any good evidence

---

[132] The literature on testimony as a source of knowledge is vast but see Lackey (2006) for an overview of some crucial themes.
[133] Also see Sterken (2015b) who makes a similar argument to explain some of the recall data I mentioned in §2.4 of Chapter 4, that indicated that people frequently recall universally quantified generalisations as generics.

for believing that wolverines are more aggressive than animal belonging to that comparison class. I want to make two clarifications before moving on to the next stage of my argument.

First, it would be wrong to say that in the cases that I have described, I have impeccable evidence for the generic belief, and no evidence whatsoever for the corresponding specific generalisations. For instance, given that the TV programme tended to be quite sensationalist, I may have reasons to doubt that there is any salient link between being a wolverine and being aggressive after all. Likewise, though I may not have any direct evidence for a more specific generalisation, I may end up with a reasonable amount of evidence in other ways. For example, imagine that I find out some additional information about different wolverine sub-species and how they developed evolutionarily. If this were the case, then I might end up believing that wolverines are sufficiently similar to one another to make it very likely that *most* of them are aggressive if some of them are.

Second, I have assumed that having good evidence is what leads us to form beliefs. However, we form beliefs without good evidence all the time, including beliefs in generalisations. It certainly isn't the case that we only believe in specific generalisations when the evidence available to us warrants it. However, it is generally accepted that there are norms governing belief and that we, to a greater or lesser extent, attempt to comply with these norms so that our beliefs are not complete stabs in the dark.[134] My claim here is that generic generalisations enable us to form justified beliefs in the absence of evidence that would be adequate for beliefs in more specific generalisations.

4.2 How is this useful?

One concern at this point would be that, even if I am right about generic beliefs and their evidential profile, these beliefs cannot be very useful after all, precisely because they are so general. What good is it to be able to believe that there is *some* generalisation out there about the kind and property in question, if we don't know much about which generalisation it is? In terms of being able to inform our actions, surely it is the fact that we know that a generalisation is for example statistical, or concerning essences, that enables us to act one way rather than another.

Even though generic beliefs are evidentially undemanding, I've argued in the previous chapter that they are inferentially powerful in the sense that they dispose those who hold them to a host of other inferences. Remember, the idea was that generic beliefs dispose those who hold them to a variety of other defeasible inferences. For one, those dispositions enable us to draw inferences about non-generic generalisations and also to expect that a given member of a kind will have the property in question.

The combination of these features, evidential undemandingness and inferential power, allows people to make the most of what little evidence we have. Even though the beliefs they end with are non-specific, they nevertheless lead those who hold them to additional, more specific beliefs, albeit defeasible ones.

---

[134] See Chignell (2018) for an overview of work on this topic.

Specifically, I argued that believing that at least one generalisation about a kind and property is true disposes us, all else being equal, to expect that several non-generic generalisations about the kind and property are true at the same time, something I called a 'co-occurrence expectation', and even expect of a particular member of the kind that they have the property in question. One thought about this might be that what I've described is a way of making reasoning errors that lead us astray from actually having evidence to forming beliefs that are not based on any good evidence. This doesn't seem useful.

It is true that generic beliefs do, as I have described them, dispose us to expect that several non-generic generalisations will be true on the basis of a generic belief, when often, they won't all be. So, generics do often lead us astray. I will consider one such type of generic belief, belief about social groups, in the next chapter.

However, whether the co-occurrence expectations of non-generic generalisations are legitimate, and hence whether we aren't merely *disposed* to further beliefs, but are *licensed* in forming them, depends on the subject matter. Co-occurrence expectations can be more or less warranted. Often, it is reasonable, if not deductively valid, to infer from the fact that there's a link between being a bus driver and suffering from slipped disks, to the belief that there's something about the work bus drivers do that disposes them to suffer slipped disks, especially if we have other well-founded beliefs, for example, that those who end up working as bus drivers don't have a significant pattern of other risk factors for slipped disks. Whether a generic belief is a *good* basis for other non-generic generalisations depends on the non-generic-generalisation in question, the kind, the property ascribed, and which prior beliefs we have. In the next chapter, I describe discuss how generic beliefs can lead us astray in the social domain.

What I have argued here is not that generic beliefs are useful in that they magically afford us justification for a variety of specific generalisations based on little evidence. If they did, that would be quite suspicious. Rather, my claim is simply that generic beliefs are non-specific generalisations that we *can* have justification for believing in based on limited evidence, even if these generic beliefs in turn lead us to forming other, sometimes less well-founded, beliefs.

The core idea in this chapter so far is this: the non-specificity of generics, the idea that they can be made true by different types of non-generic generalisations, bestows beliefs about them with a particular epistemic function. Generic beliefs allow us to make the most of limited evidence because they are both evidentially undemanding and inferentially powerful. This angle can help us better understand the cognitive significance of generic beliefs in our conceptual apparatus.

5 Conceptual role

Mark Johnston and Sarah-Jane Leslie (2012) discuss the role of generics in our conceptual apparatus. In discussing psychological theories of concepts, they point out that one feature they have in common is that 'they all make some reference to properties that are possessed by members - plural - of the target

category; they all involve focus on forms of generalization concerning the category and its properties' (2012:124).

This is not accidental. Our concept of tree, or at least the heuristics we use to determine whether something is a tree, consists in beliefs that take the form of generalisations, for instance about them being plants, having leaves, and being green. Johnston and Leslie distinguish the application conditions of a concept from the heuristics we actually use to determine whether a concept is to be applied to something. According to them, philosophers are mainly concerned with the former and psychologists with the latter.

They argue in favour of what they call the 'Generic Encoding Hypothesis':

> 'The empirical claim that the heuristics, or more generally, the criteria we use in applying terms frequently take a generic rather than a universal form' (2012:132).

They argue for this using similar data Leslie (2008, 2012) also utilises to motivate her own account, much of which I discussed in the previous chapter. Even without reviewing this sizable literature, it is easy to see why we would use generic beliefs to determine whether a given x belongs to some category rather than beliefs in universally quantified generalisations. There are many properties that are useful in determining whether something is a tree which nevertheless do not apply to all trees.

For instance, though we almost certainly use beliefs about trees having leaves and being green in our heuristics about how to determine whether something is a tree, deciduous trees don't always have leaves (e.g., in winter) and some trees are colours other than green. Conversely, not all green things or even things with leaves are trees. In addition, we categorise things not only in terms of the properties that are prevalent in members of a category, but also using properties that are typical for them, that they are disposed to have, or that they ought to have. Crucially, we often categorise things on the basis that they have some property that we know to be *associated* with the kind, even when we don't know much about the nature of the association, in just the way I described in §4. Generic beliefs then, appear to play an important role in how we represent knowledge about kinds.

Johnston and Leslie argue that generic beliefs cannot comprise the application conditions of the concepts themselves as those are compositional and generics are not. They reason that since words can combine compositionally, so should the concepts they signify. For example, the concept of a female lion should in some sense be composed of the concepts *female* and *lion*. However, the concept of a lion would plausibly involve the platitude *Lions have manes*, but the concept of a female lion wouldn't. Hence, Johnston and Leslie argue that it's not the actual application conditions of concepts that are generic, but merely the rough heuristics we use to determine whether a given x belongs to a category.[135]

---

[135] Johnson and Leslie briefly consider a cluster theory as advocated by Lewis (1972) of concepts according to which for something to belong to a category, not all, but just a cluster of platitudes needs to be fulfilled, but ultimately dismiss it. I'm not in a position to tell whether they are right in their rejection and hence don't pursue this line of argument here.

However, the generality account can explain why the concept *female lion* can include *Lions have manes* in its application conditions after all. If the generic *Lions have manes* is one of the application conditions of *lion*, then on the generality account all that means is that some generalisation about lions having manes is true. And it is: *Most male lions have manes.* However, this is not in conflict with female lions never having manes. Hence, there is no conflict between any of the application conditions of *Lion* and *Female Lion.* And so, *Female Lion* could actually include *Lions have manes* as one of its (weak) application conditions. I will not propose here that the application conditions of concepts actually are generic, much more would need to be said to do so. However, this possibility doesn't seem like something that should be dismissed out of hand and should make the generality account, and structurally similar ones, be of interest to those who want to understand what form the beliefs we use in categorising things take.

Even in a more modest form, the idea that generic beliefs form a core part of our ability to *apply* concepts is enough to prove their worth, regardless of whether the conceptual knowledge itself is generic or not. Such knowledge is an important tool for anticipating experiences. Knowledge about how to apply kind terms helps us make sense of our observations and predict future ones. Learning what dogs are and which attributes they have helps us understand patterns in experiences better, for example patterns of hearing barking when knocking on doors. It also enables us to be better able to predict the future and act accordingly. For instance, if we know that dogs bark, we can be prepared when knocking on doors.

Generic beliefs can help us understand both how we decide whether a given individual belongs to a category, by ascribing properties that are connected to a kind, however loosely, as well as how, based on that categorisation, we can form additional predictive beliefs about that individual. The explanation of defeasible inferences that generic beliefs dispose us to I gave in the previous chapter can be of use here. I argued that generic beliefs are inferentially powerful by disposing those who hold generic beliefs to making a wide range of defeasible inferences. This can explain how after we have concluded that a given x belongs to a category, for instance that Vera is a dog, this belief disposes us to make other inferences, for example that Vera barks, thus helping us to anticipate what may happen when we knock on the door of the house she lives in.

5.1 Back to Leslie once more

I argued in the previous chapter that Leslie's approach is empirically unmotivated and presents theoretical issues. Nevertheless, I believe that the approach I have presented is compatible in parts and could, pending further empirical investigations, be combined.[136] Leslie presents generics as giving voice to a default generalisation faculty. I argued that one issue with Leslie's account is that she cannot explain the (semantic) truth-conditions of generics in virtue of their cognitive role or vice versa.

[136] The semantics for generics that I presented in Chapter 3 are incompatible with Leslie's framework as I understand it because I take her to be making the specificity assumption that I argue against in Chapter 2.

The alternative presented is different: Whereas Leslie presents generic beliefs as primitive and basic in a cognitive sense tied to the dual-process approach, the generality account suggests that they are basic in an *epistemic* sense, that is to say not tied to a dual process approach. It could still turn out that generic beliefs have a particular cognitive primitiveness because of their epistemic features or vice versa. For example, we might think that generic beliefs are acquired early because they are so epistemically basic that they are a way for children to form successful beliefs about the world based on little evidence. Or it might be that generics are prevalent in parental speech just because parents use them to communicate heuristics by which to determine whether a given object belongs to a category, or even spell out a concept's application conditions. This would be an extension of the view that I have presented here and is, at this point, mere speculation. However, it is worth pointing out that, although I have presented the generality account as an alternative to Leslie's framework, aspects of it are in fact compatible with some of the empirical claims Leslie makes and could be combined.

The argument I have put forward in this chapter is not only compatible with Leslie's view, but also other functions that might be central to understanding generics. For example, non-specificity might also just allow us to form beliefs and communicate them in a way that optimises cognitive bandwidth if it easier to form, store, or otherwise process beliefs in non-specific generalisations than the more specific generalisations that I've argued make them true. I haven't explored this argument here because I take its plausibility to depend on the outcome of experimental work that is yet to be conducted. Nevertheless, it would be an interesting line of argument to pursue as it could afford an alternative understanding of generics as being 'default' in a way that isn't as closely tied to a dual process approach as Leslie's account.

So far, everything discussed in this chapter has related to generic beliefs. In the last section, I will return to generic sentences.

6 In speech

Given the usefulness of non-specific but stable contents in cognition that I have argued for in the previous section, it should come as no surprise that I also think it is useful for us to be able to verbally express such content. This is the narrow linguistic function Cappelen talks about as the only function of expressions. Generic sentences are useful in this sense because they enable us to communicate our beliefs to others. In this section, I explain how having non-specific generalisation be *both* the semantic content of generics and the objects of generic belief can explain how we are able to meet norms of assertion for generics through generic knowledge. I then end by considering whether this makes generics a type of hedging speech.

6.1 Norms of assertion

Generics are declarative sentences, not interrogatives or imperatives. Given that declaratives are conventionally used to make assertions, my assumption is that utterances of generics are assertions.

However, assertions of generics are, like all other assertions, governed by norms of assertion. The idea that assertions of generics are governed by specific norms is based on the general idea that assertions are speech acts that claim that something is the case. Consequently, there can be assertions that are in some sense proper, and assertions that are improper. An assertion is improper if the proposition asserted is false, but it also seems improper if the asserter is just guessing.

Norms of assertion can state the conditions under which an assertion is proper in this sense. Note that assertions might be improper in all sorts of ways, some having nothing to with what an assertion should be like as an assertion. For example, an assertion might state something that actually is the case, but nevertheless be improper in a different sense, because it is rude. This is not the sense of propriety that norms of assertion address, rather, they are about the sense in which assertions can be proper/improper *as assertions*.

The most famous norm of assertion is Timothy Williamson's knowledge rule (1996). According to Williamson, one can properly assert P only if one knows that P. The knowledge rule can capture both what's wrong with asserting something that's false, as well as what's wrong with asserting something that is true, but for which one has little or no evidence. As knowledge is assumed to be factive, asserting something false is improper because one cannot know it. However, something's being true doesn't guarantee that we have evidence that this is the case, especially not evidence that meets the threshold for knowledge, and so one can assert that P improperly even when P is true if one doesn't have good enough evidence that this is so.

Williamson's knowledge norm has been very popular, however others who also believe that the knowledge norm includes an epistemic component disagree that that component requires anything as strong as knowledge. For example, Lackey (2007) and Kvanvig (2009) focus on justification, and Maitra and Weatherson (2010) emphasise the importance of one's attitude being evidence responsive. I will be using the knowledge norm as an example of a norm of assertion just because it is popular; however, because I have framed the epistemic function of generics in terms of evidence, what I have to say here will apply to any norm of assertion that involves an epistemic requirement.

In the first section of this chapter, I argued that generic beliefs are more easily justified than beliefs in non-generic generalisations because they can be made true by a much wider class of evidence. If this is so, then it can explain why it is easy for speakers to abide by the relevant norms of assertion for generics. The relevant norms require that to assert a generic of the form *Fs are G*, a speaker must know that Fs are G. I've argued that it is easy for speakers to know that Fs are G, at least in comparison to knowing the more specific generalisations expressed by overtly quantified generalisations. Consequently, it is easy for them to comply with the norms for asserting generics.

Contrast this with overtly quantified generalisations. In cases in which a speaker doesn't know, for example whether most wolverines are aggressive or whether wolverines are more aggressive than martens, they are unable to assert the corresponding overt generalisations without violating norms of assertion.

However, as long as they know that *some* salient generalisation is true, they can properly assert *Wolverines are aggressive*.

Of course, it isn't surprising that the contents of generic sentences, if known, enable speakers to assert the generic sentences in question. The same goes for any sentence and its contents. My point is rather to emphasise that thinking of generics as having distinctive contents and of generic beliefs as attitudes towards *those same contents* can give us a straightforward explanation of how generics can be used to give voice to our generic beliefs in contrast to accounts like Leslie, which do not preserve such a direct link between the two.

In contrast, accounts that posit non-specificity as a result of contextualism or semantic incompleteness have no distinctive role for generic beliefs. Hence, if they agree that at least some generic beliefs are evidentially undemanding, they must understand non-specific generic beliefs as attitudes to sets containing several, or no propositions. Propositional nihilists seem to me to have little hope of giving a straightforward explanation of what non-specific generic beliefs consist in. Pluralists on the other hand must explain how we can take propositional attitudes of belief to sets containing *more* propositions, seemingly on the basis of *less* evidence. This seems difficult to me. Likewise, both nihilists and pluralists must give an additional account of how norms of assertion govern assertions that convey several, or no propositions. I hope to have shown that the non-specificity account can provide a simpler, and more attractive explanation.

6.2 Hedging

One question we might have in mind with regards to communicating generalisations is this: if we have limited evidence for the generalisations we want to express, why not just hedge overtly quantified generalisations? One way for speakers who have access to limited evidence to avoid violating norms of assertions is to minimise their conversational commitment by hedging what they say.

Matthew Benton and Peter van Elswyck (2018) characterise hedges as linguistic constructions that lessen conversational commitment. They use phrases like 'probably' or 'I think' as their paradigmatic cases of hedging. According to Benton and van Elswyck, these expressions are added onto sentences, for which speakers would otherwise fail to meet the epistemic requirements required by the norm of assertion. For example, I don't know that it's going to rain tomorrow, although I think it's likely. I would not meet the knowledge norm for asserting *It's going to rain tomorrow*. However, I can hedge my assertion by adding 'probably'. According to Benton and van Elswyck, adding hedges like *probably* signal that speakers lack the knowledge required for the unhedged assertions. By signalling that their pre-hedged assertion doesn't meet the knowledge norm, speakers advertise that their assertion is less certain than a non-hedged

counterpart. According to Benton and Elswyck, this is how speakers lessen conversational commitments and thereby comply with norms of assertion.[137]

So, why can't speakers just hedge overtly quantified generalisations? They can, and often do. However, being able to *assert* generalisations using generics has benefits derivative of the epistemic function I described in the previous sections. In asserting, we represent ourselves as knowing something, which, other things being equal, gives our interlocutors reason to believe what we say. Assertions are thus a means to communicate knowledge. And, as I've argued in §3 and §4, we can have evidence that provides us with very good justification, even to the point of giving us knowledge, for a given generic belief even if that evidence is not particularly good evidence for believing a given non-generic generalisation. Even though I might not know in virtue of what it is true that wolverines are aggressive, I may nevertheless have excellent justification for believing that at least one non-generalisation about them is true. Generics are what makes it possible to communicate this knowledge through assertion.

Elin McCready (2018) distinguishes between force and content hedges. Force hedges include those Benton and van Elswyck use as examples, like 'perhaps' and 'I think'. These constructions, when added to full sentences, weaken the force of the utterance to the point where she classifies them as no longer assertions. However, McCready places just as much emphasis on another type of hedge, *content hedges*. These are expressions like 'kind of' or epistemic modals like 'might', which don't weaken the force of the sentence in which they appear, but its content, making it harder to falsify. According to her, these content hedges appear in full assertions, as their *force* is unaffected by the hedge.

Generic assertions are content hedges. Their content is non-specific. This non-specificity weakens the semantic, though not the pragmatic, content of the generic asserted because it is non-specific with regard to which generalisation about a target kind and property is true and requires only that at least one of them is. As a result, speakers can easily know the content, easily enough to properly assert the generic in question even when they would not be able to properly assert any one non-generic generalisation. This makes generics part of a class of content hedges, along with expressions like epistemic modals, which enable speaker to communicate via assertion in imperfect epistemic circumstances.

7 Conclusion

In this chapter, I have used the generality account to explore a set of interrelated functions of genericity. I argued that because the content of generics is non-specific, generics are easy to believe and know, which in turn makes them easy to assert. Given the central role generic beliefs play in our ability to categorise, being able to form such beliefs and give others reason to do likewise is extraordinarily useful. Insofar as

[137] I'm unsure about how flagging that one is not complying with a norm makes it such that one complies with it but will not pursue that question here.

the generality account can explain this usefulness and hence illuminate the role of genericity in cognition and communication, I take that to speak in its favour.

Chapter 6: Generics and Stereotyping

1 Introduction

Up to this point in the thesis, I haven't restricted my focus to generics about particular kinds or properties. That will change in this chapter.

Consider the following:

(1) Blondes are stupid

(2) British people have bad teeth

(3) Bus drivers are grumpy

(1)-(3) are generics that express stereotypes about blondes, British people, and bus drivers.[138] A lot of recent interest in generics has been based on the thought that there is some particularly close connection between generics and stereotyping.[139] In this chapter, I want to offer a new diagnosis of the connection that makes use of the account of generic beliefs I have developed in Chapters 4 and 5.

I'll start by introducing generics and stereotyping separately (§2) and discuss some initial reasons for thinking that there is a close link between them in §3. I then offer my own diagnosis (§4), namely that stereotypes consist of generic beliefs alongside non-propositional elements. In §5, I apply the considerations from the last chapter to the case of stereotyping. The generic contents of stereotypes enable speakers to form evidentially undemanding, but inferentially powerful beliefs about social groups, which they can express in a manner that preserves plausible deniability. In §6, I consider an objection, namely that this account predicts that stereotypes are true. I respond by arguing that stereotypes dispose us to characteristic false inference patterns about social groups (§7) and that in verbally expressing them, speakers negligently propagate falsehoods. I end by considering how my account can inform responses to stereotyping (§8).

2 What is stereotyping?

I'm going to be interested in stereotyping as an activity that causally involves stereotypes. I will first clarify what stereotypes are and then explain how they feature causally in stereotyping. Stereotypes are mental states associating groups with properties.[140] Katherine Puddifoot defines stereotypes as 'mental states that

[138] (1)-(3) are not particularly bad examples of stereotypes. For reasons I explain in §7.2 of this chapter, I don't want to repeat more harmful ones unnecessarily.

[139] For example, see Anderson et al. (2013), Beeghly (2015), Haslanger (2011, 2014), Leslie (2014, 2015, 2017), O'Donnell (2017), Ritchie (2019), Saul (2017), Lemeire (2020).

[140] See for example in Puddifoot (2017), Beeghly (2015), Blum (2004), Jussim et al.(2015).

associate members of certain social groups more strongly than others with particular attributes, in virtue of their perceived social group membership' (2017:139).

This definition is inclusive in one sense but narrow in another. It's inclusive in that it doesn't restrict the types of properties that can be represented in stereotypes to, for example, stigmatised properties. According to Puddifoot's definition, mental states with the content of (4) and (5) are stereotypes, while those with the contents of (6)-(8) are not.

(4) Brazilians live in Brazil

(5) Children get chicken pox

(6) Ottawa is boring

(7) Blunt knives are dangerous

(8) Pebbledash houses are depressing

Given that there isn't anything wrong with living in Brazil and chicken pox though unpleasant for those afflicted isn't stigmatised, we might want to take issue with saying that (4) and (5) are stereotypes. In contrast, even though (6), (7), and (8) aren't about social groups, they ascribe negative properties and generalisations like them are often called 'stereotypes'.

Whether these should or shouldn't count as stereotypes depends on what one thinks the notion of a stereotype does or should do. Those who want to emphasise the continuity between stereotyping and benign categorisation processes are happy to embrace an inclusive notion of stereotypes. Those who see stereotyping as the malignant outgrowth of those processes favour a more restrictive notion.

Though, as will become clear, I see the malignant element as arising out of a general function of generic beliefs, which also underlie the benign uses, I am primarily interested in stereotyping of groups of people in which stigmatised, or otherwise negative properties are ascribed to them. For that reason, I will focus on those stereotypes that meet Puddifoot's definition and remain neutral on whether generalisations like (4) - (8) should or shouldn't be called 'stereotypes'.

Stereotyping can comprise different types of activities. One important thing to note is that stereotyping can, but need not, involve verbal expression. Take the following example. I hold a stereotype about fans of snooker player Ronnie O'Sullivan, which is that they are prone to heckling. As a result, I might expect a given Ronnie O'Sullivan fan to be prone to heckling and treat them in a hostile manner. This seems to be an example of stereotyping, regardless of whether I express the stereotypes I hold verbally.

Erin Beeghly (2015) distinguishes between mere stereotype activation and various ways in which stereotypes can further feature in judgments. Stereotype activation occurs when a stereotype is triggered in some way. The stereotype can then bring about beliefs, justify behaviour, or be expressed in speech. First, the generalisation can play a causal role in forming beliefs. For example, my belief that Ronnie

O'Sullivan fans are prone to heckling, together with the belief that Briony is a Ronnie O'Sullivan fan, can cause my belief that Briony is prone to heckling.

Second, stereotypes can feature as reasons for a course of action. For example, if I were to attend a snooker tournament, I might try to find a seat far away from Briony because I believe Ronnie O'Sullivan fans like her to be prone to heckling. Third, the stereotype can be communicated linguistically. For example, I might explain to my friend that I would like to find seating elsewhere by pointing to Briony and saying, 'Ronnie O'Sullivan fans are hecklers'.

Often, several of these types of stereotyping occur simultaneously or in succession, but they need not. In the example, my stereotype about Ronnie O'Sullivan fans led to my forming the belief that Briony was prone to heckling, which in turn gave me a reason to find another seat, something I justified by expressing my stereotype verbally. However, many cases of stereotyping involve some, but not all types. For example, I may form the beliefs about Briony and avoid her, but not say anything. This is why it's important not to focus on verbal stereotyping only.

3 What's the connection between generics and stereotyping?

There are a few different considerations that seem to point towards the idea that there is a particularly close connection between generics and stereotyping. The first is that speakers often use generics rather than overtly quantified generalisations to express stereotypes they hold. For example, instead of saying 'All blondes are stupid', they will say 'Blondes are stupid'. Many strengthen this claim to say that stereotypes are usually, or typically, expressed using generics. For instance, Lawrence Blum calls generics 'the typical form of a stereotypical generalization' (2004:259), while Beeghly says that 'Stereotypes expressed in speech tend to take the form of what linguists call generics' (2015:676).

I agree that the use of universally quantified generalisations in expressing stereotypes is rare, although it does occur. One thing to point out is that Blum and Beeghly contrast the use of generics with that of universally quantified generalisations. However, universally quantified generalisations are not the only kind of overtly quantified generalisation. While I agree that speakers infrequently use 'All blondes are stupid', it's not obvious to me that generics are much more prevalent than generalisations like 'Most/Many blondes' are stupid or ones like 'Blondes are usually/often/generally stupid'.

Anecdotally, my impression is that generics are more prevalent, but I am not aware of any empirical data on the prevalence of generics in expressing stereotypes as opposed to various overtly quantified generalisations. This is an unfortunate gap in the literature. My argument will provide some reason to think that we should expect generics to be prevalent in stereotyping speech, and so should be taken as a hypothesis that can be tested empirically.

Another intriguing piece of evidence is that, as Blum points out, stereotypes are typically 'held in a manner that renders them largely, though not entirely immune to counter-evidence' (2004:251). This is

particularly interesting given that one aspect of generics that is often thought to distinguish them from universally quantified generalisations is that they tolerate exceptions. The fact that penguins can't fly makes *All birds fly* false, but *Birds fly* remains true. This behaviour mirrors the ease with which stereotypes can survive what seem like exceptions. Not very many who hold stereotypes about blondes being stupid, believe that all blondes are stupid. As a result, they don't take the odd non-stupid blonde as a reason to change their mind.

Here's a straightforward explanation for these data points: the content of stereotypes is generic. Beeghly considers whether stereotypes might be a subset of generics but points out that they can't because generics are linguistic, and stereotypes are not.[141] However, I think that there is a view in the vicinity that is worth exploring in more detail. It's true that generics are sentences and stereotyping need not be linguistic. However, the content of generics, the propositions they express, might nevertheless have a distinctive generic character that could be the object of propositional attitudes involved in stereotyping. I present such a view in the coming sections.

4 The proposal: stereotypes involve generic beliefs

Here is the basic idea. Generics are linguistic items that express propositions, generic propositions. We can take different attitudes towards such propositions, we can hope for them, dread them, expect them, or we can believe them. In cases where we believe a generic proposition of a certain type, very roughly concerning a social group, we hold a stereotype about that group. When this belief causes other beliefs to be formed, serves as a reason for action, or is expressed verbally, we are stereotyping.

I will say much more about the nature of these generic beliefs shortly, but I first want to consider an initial objection. The objection concerns thinking of stereotypes as belief-like at all. As beliefs are propositional attitudes, they are attitudes taken towards propositions. However, many believe at least some of the content of stereotypes to be non-propositional. If that were so, at least some of the content of a stereotype could not be the kind of thing that is or isn't believed. Hence, stereotypes couldn't just be kinds of belief.

For instance, Eliot Smith and Jamie DeCoster (2000) argue that our associations with social groups include things like emotions and images that are not propositional. Ella Whiteley (forthcoming) argues that stereotypes involve patterns of attention that influence which characteristics of individuals are salient to us. These patterns of attention are non-propositional and therefore cannot be part of beliefs. I am sympathetic to their arguments and hence do not want to try to accommodate these non-propositional components within a standard model of belief. Instead, I want to argue that although stereotypes don't solely consist of generic beliefs, they do have a belief component.

---

[141] Beeghly (2015:677)

What I have said so far is independent of any particular account of generics. Applied to the case of stereotypes, the view is as follows. Stereotypes are mental states that have a belief component, alongside other, non-propositional components like emotions, images, or patterns of attention. The belief component takes the form of a generic belief.[142] On this view, stereotypes are subsets of generic beliefs about social groups.

Stereotyping is the featuring of stereotypes in cognitive processes, communication, or as justifications for actions. Generic beliefs, including stereotypes about social groups, can cause other beliefs, act as justification for actions, and be communicated, often in the form of generics. Let's take the stereotype of blondes being stupid. This stereotype has a belief component and other, non-propositional components. The belief component is the belief that blondes are stupid. This belief is an attitude taken to a proposition expressed by the generic *Blondes are stupid*. The belief that blondes are stupid can cause other beliefs, for instance that Josefine, qua blonde, is stupid; cause actions, to not hire her; or can be expressed by saying 'blondes are stupid'.

5 The inner workings of stereotyping

In this section I want to draw attention to the ways in which accounts of generics, and in particular the generality account, can explain aspects of stereotyping. The features of generics that enable these processes are not specific to the contents of the generics in question, and I explored them in the previous chapter. However, I will argue that the hedging mechanism is particularly useful to speakers for verbally expressing stereotypes as it enables them to communicate stereotypes while maintaining plausible deniability.

5.1 Generic beliefs

Generic beliefs about social groups are useful to us in the same way generic beliefs about anything are. Generic propositions enable people to form non-specific beliefs about social groups, which are easily justified and highly useful in categorisation. Given the importance we attach, rightly or wrongly, to categorising people as belonging to social groups, generic beliefs are a central tool in identifying group membership as predicting what individuals are likely to be like based on the groups to which they belong.

Of course, the same mechanism that makes stereotypes useful in this way, is also a core reason for why they are harmful. Because stereotypes dispose those who accept them to make defeasible inferences about particular individuals that belong to a category having some property, they can form false, or otherwise harmful beliefs about individuals and the groups they belong to.

---

[142] For some experimental data about why the belief components in question are generic rather than statistical see Hammond & Cimpian (2017)

These concerns are borne out by empirical evidence. The work in question typically focuses on the connection between generics and essentialising. The psychological phenomenon of essentialising occurs when people believe members of social and natural kinds to have a fundamental nature, which is thought to determine their more superficial properties.[143] Several studies have indicated a link between generics and essentialising. For instance, studies suggest that receiving information about kinds in the form of generics leads us to essentialise the kinds. This is so, even in research that made use of generics about novel social and natural kinds about which we have limited background beliefs.[144] [145]

This link can be explained by the account of generic belief I proposed in Chapter 4, and specifically, the tendency of generic beliefs to give rise to co-occurrence expectations. Because we have an underlying assumption that if one type of generalisation about a member of a kind is true, others are too, we, barring incompatibilities with other background beliefs, come to expect both that there is a deep reason for prevalences as well as that prevalences occur in the first place.[146] Essentialising is one instance of this broader inference pattern and happens when the co-occurrence expectations lead us to infer from there being *some* connection between a kind and property, that that connection is grounded in underlying, stable properties shared by members of the kind. This phenomenon is not specific to stereotyping, but it is especially troubling as research indicates that the most highly essentialised groups also experience the worst prejudice.[147]

A recent paper by McKeever and Sterken (2021) describes four of distinctive, flawed inference patterns that generic stereotypes give rise to: amplification, degradation, shifting, and entrenchment. I want to focus on two in particular and suggest that these are instances of the inferential pattern described above. *Amplification* occurs when individuals accept generic claims based on weak evidence, but then take such evidence to warrant strong claims. For example, individuals form the belief that blondes are stupid on the basis of relatively weak evidence, but then take that belief to warrant stronger statistical, modal, and normative claims. *Shifting* occurs when individuals accept a generic based one a non-generic generalisation of a particular flavour, but later take that same generic to support a different non-generic generalisation. These inference types can be seen as characteristic patterns of co-occurrence expectations of the kind I described in the previous chapter.

---

[143] See Leslie (2013) for a discussion on the relationship between the psychological phenomenon of essentialising and philosophical essentialism. I'm concerned with the psychological notion.

[144] For example, see Cimpian, Meltzer, and Markman (2011), Cimpian and Erickson (2012), Cimpian and Markman (2008).

[145] Research by Hoicka et al (2018) suggests that essentialising may not be the purview of generic generalisations only but can also be caused by hearing non-generic generalisations. This is compatible with the explanation of stereotyping I have given in this section as well as for example the argument given by Munton (see fn 142), which describe a similar mechanism in the case of non-generic generalisations.

[146] See Munton (2019) on a similar argument for overtly quantified generalisations. As I noted in Chapter 4, arguing that generic beliefs dispose us to co-occurrence expectations doesn't preclude other beliefs, including beliefs in overtly quantified generalisations, from doing the same.

[147] Haslam, Rothschild, and Ernst (2002, 2000)

In both cases, generic beliefs are formed (based on more or less specific evidence) which, barring incompatibilities with background beliefs, in turn give rise to co-occurrence expectations and hence, beliefs in more specific generalisations. Our prior beliefs about social kinds and the nature of the stigmatising properties that are typically ascribed in stereotypes interact with cognitive biases to give rise to the distinctive inference patterns McKeever and Sterken describe. The same mechanism, inferential power on the basis of limited evidence, that, I argued in Chapter 5, can make generics useful has a sinister social effect.

5.2 Asserting generics

One noted feature of stereotyping that involves generic sentences that speakers appear to be able to use them, but evade responsibility for specific, offensive generalisations when challenged.[148] For example, in asserting 'Blondes are stupid, the speaker can, when challenged, avoid having to take responsibility for the claim that most blondes are stupid, or that blondes are stupid in virtue of some sort of shared nature.

The generality account has a straightforward explanation of this. Using generics with non-specific contents enables speakers to retain plausible deniability because their utterances, at least semantically, don't express any one specific generalisation for which they could be held accountable. Of course, the speaker doesn't have to take responsibility for the claim that most blondes are stupid, or that they are stupid because of some sort of shared nature – simply because they didn't express any of those claims.

In uttering a generic, speakers don't have to assume responsibility for any of the non-generic generalisation that could make true the generic, only that at least one of them is true. This is a general feature of generic assertions, but it is especially useful for speakers when the content communicated is controversial. To be clear, my claim is not that speakers, in general, choose to employ generics for this reason. Sometimes this is the case, but often we use generics simply to express the non-specific generalisations we have in mind.

5.2.1 Testing alignment

Using generics help us communicate in ways that are not purely aimed at sharing information, at least not about the subject matter. I consider one such way, using generics to test alignment with one's interlocutors, here.

We don't just say things to provide information to others, but also to gather it from them. We often say things to see how others react. This may either be as a way to get confirmation or to gain additional information. Our interlocutors' reactions can be useful as a source of information with which we can scrutinise our own beliefs, but also to inform our future interactions with them. We might say something to see if our beliefs are widely held, or to get others to give us information about the topic that we don't

[148] As observed by Anderson et al (2012) and Lemeire (2020)

yet have access to. Generics are useful for such purposes, in particular in cases in which the content communicated is controversial. Speakers can *test the waters* by uttering generics to ascertain the extent to which their interlocutors share their general views, but without making themselves vulnerable to challenge. Let me demonstrate this using an example.

Edward believes biological differences make women more emotional than men and that, because of this, they are disposed to be more emotional than men. Edward and his colleague Rob witness a female colleague be emotional. Edward says 'Women are emotional' to Rob. If Rob takes issue with this and challenges him, Edward can defend his assertion by pointing out that he didn't *say* any specific, more controversial claims. Often speakers in such situations will point out that their utterance is made true by a type of generalisation less likely to cause offence. By having fallback truthmakers, speakers can utter generics confident that they will be made true *somehow* and thereby not have to retract their assertions or face the social cost of being known to hold controversial views. However, if Rob does not take issue with this, or even volunteers further information that indicates that he thinks the assertion is made true by a fact about women's natures, Edward can feel safe in knowing that his and Rob's views on women align.

This example has quite an intentional flavour, as Edward is using non-specific language to see whether he would be safe to express his more specific, but also more contentious beliefs. However, testing alignment need not be so intentional. For example, Rob might have some general sense that women are more emotional than men, but not settled opinion about in virtue of what this is so. He could also say 'Women are emotional', both to see to what extent his general view is shared, but also to elicit more information that would allow him to make up his mind. Generics thus enable speakers to test alignment or elicit further information with others in a way that minimises the risk of being held accountable or of having to profess one's ignorance.

5.3 Alternative explanations

The case of stereotyping provides another useful test case for some of the accounts of generics I have discussed in this thesis. For the generality account in its invariant form, stereotyping is an instance of a more general phenomenon, namely the taking of attitudes to non-specific generalisations, in this case about social groups. What about those that posit semantic incompleteness or contextualism in generics, 'the meta-semanticists'?

The meta-semanticist qua meta-semanticist isn't particularly interested in the mental states generics give voice to, but more so the way they can be expressed in language. I noted this in the previous chapter, but it bears pointing out in the case of stereotyping in particular. If there are no distinctive generic propositions generics express, then there are no distinctive non-specific objects of belief that could play the role of stereotypes, as I have argued is the case in this chapter.

Where the meta-semanticist can contribute is in explaining verbal stereotyping using generic sentences. In particular, they can offer different accounts of how speakers who stereotype using generics are able to

maintain plausible deniability. Those who posit semantic incompleteness have a straightforward explanation that is similar to the one I gave above. Speakers who stereotype by using generics can evade responsibility for any one non-generic generalisation that would clearly be offensive because the generic token doesn't itself express any such generalisation, at least not semantically. Instead, the controversial content is conveyed through impliciture, which, like other pragmatic content, can be disavowed more easily than semantic content.

The contextualist on the other hand has two options.[149] Either speakers who stereotype verbally semantically express non-generic generalisations, or the generalisations they express are non-specific to varying degrees. In the case of the former, speakers who for instance have the relevant specific intentions, like Edward, may utter a generic which expresses an offensive non-generic generalisation, but they can deny this by disputing the relevant intentions when challenged. This explanation paints the evading of responsibility as an intentional practice.

Alternatively, contextualists may say that cases of verbal stereotyping involve the communication of generalisations that are non-specific to varying extents and embrace propositional nihilism (in the way described above) or propositional pluralism. The pluralist will then have to explain how it is that generics can be used by speakers to preserve plausible deniability via a theory of how speakers are held to account for the sets of propositions their utterances express. This might be difficult. For example, Viebahn (2019) argues that to assert a sentence that expresses several propositions a speaker must meet epistemic requirements for each of the propositions expressed, which would predict that speakers shouldn't be able to retain plausible deniability as easily as others and I have suggested. [150]

6 So, is it a stereotype because it's true?

In this section I want to discuss a possible objection to my argument so far. The objection is a special case of the more general overgeneration worries the generality account faces, and which I discussed in chapter 3. It's a special case, in that it only applies to a subset of generics, namely ones about social groups. However, it's also a special case in the sense that it's especially bad. I have argued that stereotypes involve generic beliefs. The problem arises due to my account predicting that because generic beliefs and linguistic tokens are by default non-specific, they can relatively easily be made true. Specifically, as long as one of the contextually salient relations actually obtains, the generic in question is true. I've argued that this non-specificity is key to the success of generic generalisations in stereotyping. It explains why generic beliefs are evidentially undemanding and why generics enable the communication of stereotypes while

[149] I don't mean to suggest that contextualists must argue that verbal stereotyping is always of one kind or the other, only that for each individual case these are the explanatory strategies as I see them.
[150] For another view, see von Fintel and Gillies (2011). Their view is complicated slightly because they don't think that sentences of the kind they are interested in straightforwardly express several propositions, but merely 'put them into play' (2011:119). However, they provide suggestions for norms governing how this is done, and also how such utterances are taken up or denied.

preserving plausible deniability for speakers. However, the same non-specificity seems to also entail that many generic beliefs are true. Sally Haslanger (2011) discusses the following example:

(9) Women are submissive

Haslanger argues that because of the way women in patriarchal societies are raised, many of them, perhaps most, do in fact behave submissively. However, if this is so, then according to the view I have presented (9) is plausibly true because at least one generalisation about women being submissive is true. The worry then is that far from equipping us with the tools to challenge instances of stereotyping, the account I have proposed affirms them.

There are two parts to my response: First, I consider a contextualist response. Might the truth-conditions of generics that express stereotypes be systematically shaped by essentialist biases in such a manner to make them more specific and thus less likely to be true? I argue that this response is misguided and that we should instead focus on the inferential patterns stereotypes dispose us to. This can explain how generics used to express stereotypes are misleading, independently of whether they are true.

6.1 Might they be false?

The contextualist has a possible response to the worry that expressions of stereotypes might turn out to be true. They could argue that the contexts in which (9) is uttered share features that make it such that (9) is in fact relatively specific, and because of this, false.

For the contextualist, if the contextual features make salient only a certain type of non-generic generalisation, then generic tokens are specific. It's only when the contextual features are unrestrictive enough to have several non-generic generalisations be salient that the generic tokens are non-specific to varying degrees. These contextual features can be influenced by conversational moves, but also by long-standing background beliefs that are shared between participants. This is relevant in the case of stereotyping generics because we have entrenched background beliefs about frequently stereotyped groups like gender or race groups.

The idea then is that given that these background beliefs are widely shared, even contexts in which there hasn't been any explicit mention of, say, women's nature are such that essentialist association relations are the most salient ones. If that's right, then generic tokens asserted in such contexts would express more specific generalisations about women, unless the speaker explicitly disavows the intention to convey such contents. If the essences that would be required to underwrite such generic tokens don't exist, the tokens could then be said to be straightforwardly false. This response might address some worries about whether the generality account predicts that stereotypes are true. Nevertheless, I think it has its limits.

For one, unless we want to say that contexts in which highly essentialised groups are discussed are such that the only generalisations that can be salient are ones concerning the group's shared nature, we will not be able to rule out the possibility that some ascriptions of properties that are prevalent, or which

members of the group are disposed to have will be true. It just seems implausible to me that contexts in which we are discussing what most women are like or what women can be like, would not as a result have prevalence or dispositional generalisations be salient enough to have a generic token be made true by them. But if that is so, then some expressions of stereotypes about members of some gender, ethnic, sexual orientation groups etc. will, in particular contexts, be true.

The sting might be taken out somewhat by the fact that those contexts would be ones in which it was clear that the investigative aims of the conversations were about the prevalence of, or disposition to possess, certain properties, and not about whether members of those groups have the superficial properties in question in virtue of some shared nature. I think that this response by itself is still unsatisfactory, and it is made even less satisfactory by another consideration. In the previous sections, I argued that it is the non-specificity of generics that enables successful stereotyping. Because the contents of generics are by default non-specific, they enable the forming of evidentially undemanding beliefs and the communication of those beliefs in a way that limits the speaker's conversational risks. If this is right at all, then it should also extend to stereotypes about highly essentialised groups. It shouldn't be harder to stereotype those groups most frequently stereotyped. The problem is that the exact reason that renders generics so useful as stereotypes, their non- specificity, is also what can make them true. In the next section, I will move away from focusing on the truth-conditions of generic tokens and instead look at their pragmatic effects.

7 A wrong in stereotyping

7.1 Pragmatic effects of verbal stereotyping

Even if we don't want to say that stereotyping generics always strictly express falsehoods, we may wish to say that they nevertheless communicate them and can be challenged on that basis (more on whether and how we should challenge them shortly). In chapters 3 and 4, I explored the idea that generics have distinctive pragmatic effects. Drawing on the work of Haslanger (2011, 2014), I argued that generics convey implicatures based on cognitive biases, background assumptions, and relevance expectations about the kind and property, respectively. This picture can give us a sense of what might be wrong with utterances of generics, independently of what they strictly express, in terms of the implicatures they convey.

(10) Boys don't cry[151]

What's wrong with saying 'Boys don't cry' is not what the utterance strictly expresses, but what it implicates - namely that they shouldn't. I argued that the generality account can supplement this picture of implicatures by explaining why they are conveyed. Generic tokens are often too non-specific to be in

[151] Example given by Leslie (2015)

conflict with the background beliefs that generate the implicatures in question. It's because (10) doesn't specify the nature of the link between being a boy and crying that our background assumptions about differences between boys and girls, and what those differences are based on, can fill in the gap to implicate that boys shouldn't cry. This is especially so, when (10) is uttered in a context in which a normative claim would be relevant and isn't ruled out by other, prior, beliefs.

This picture of implicatures could also explain why generics cause those who believe them to essentialise. At least some of the implicatures, specifically those that are generated by essentialist background assumptions, articulate exactly those assumptions involved in essentialising, namely that members of a kind have their superficial properties in virtue of some shared, underlying characteristics.

7.2 Misleading with generics

This explanation could substantiate the claim that generic assertions that express stereotypes are misleading. If uttering a generic disposes interlocutors to false beliefs in virtue of its implicatures, then the utterance misleads, even if it itself is true.[152] One might worry that to mislead requires the intention to deceive. Given that I've argued that non-specific generic contents facilitate the forming of non-specific beliefs, which are then expressed using generics, it doesn't seem as though speakers would, in most instances, have specific generalisations in mind that they want to mislead their interlocutors about. I agree that speakers don't need to have such intentions, however, that doesn't stand in the way of characterising their utterances as misleading.

Jennifer Saul (2018) has coined the term 'negligent falsehoods' for exactly such utterances. Negligent falsehoods are falsehoods 'propagated without sufficient attention to ascertaining the truth of the matter.' (2018:249) According to Saul, this can occur when speakers' utterances are true, but nevertheless cause interlocutors to believe falsehoods. This is a useful characterisation of what speakers do when they express stereotypes using generics. Speakers assert generic tokens that, even if they are non-specific enough to be true, dispose their interlocutors or themselves to believe falsehoods about members of the group the generic is about. Given the harm the falsehoods cause, uttering such generics is negligent at best. So even if we can't say that generics used to express stereotypes are always false, we can still challenge them on the basis that they are misleading.

7.3 How to respond

One prominent strategy in the literature on oppressive speech, of which stereotyping generics are often cited an example, takes the form of *counterspeech*. Several authors that take this strategy have offered instruction on how speech whose harm comes about as a result of pragmatically communicated content

[152] I focus here on generic tokens that implicate *false* beliefs about stereotyped groups. However, harmful beliefs may not always be false. Munton (2019) argues that true non-generic generalisations may nevertheless be harmful as a result of a process that is very similar to the co-occurrence expectations I have described.

should be responded to. For example, Rae Langton (2018) argues that we can stop content communicated implicitly through presupposition by challenging the authority of those who presuppose it and thereby prevent the presupposition from entering the common ground of the conversation through accommodation. Marina Sbisà (1999) has argued that presupposed, or otherwise implicitly conveyed content should be made explicit through what she calls 'explicitation procedures' in order to make it possible to challenge them and hold speakers accountable.

Sbisà's response in particular would seem to be appropriate in the case of generics. If what makes it hard to hold those who stereotype using generics to account is that the offensive content is communicated implicitly, then what better response than to force them to make what they say explicit? The idea would be that when someone uses a generic about social groups that their interlocutors have reasons to believe conveys a stereotype, those interlocutors should get the speaker to be specific about which generalisation they intend to convey so that they can be held to account.

Alternatively, Haslanger (2011) suggests using meta-linguistic negation to challenge the pragmatically conveyed content of stereotyping generics.[153] Her idea is that we can simply deny generics, not on the basis that what they semantically express is false, but on the basis that their pragmatic content is. So, for example, was someone to utter (9), *Women are submissive*, Haslanger recommends rejecting their utterance on the basis that it coveys, via implicature, the false claim that women are, by their natures, submissive.[154]

An argument by Mary Kate McGowan (2009) should give us pause for thought. McGowan argues that the counter-speech strategy is not sufficiently attentive to real-life conversational dynamics. She argues that oppressive speech has a 'sticky' quality to it that makes it hard to undo using more speech. Once a conversational norm is enacted through oppressive speech, trying to undo that norm is like trying to 'unring a bell' (2019:403). Robert Simson (2013) has generalised McGowan's argument by tracing the stickiness of oppressive speech to a general mechanism, namely that it is much easier to prevent something becoming salient, than it is to make something already salient unsalient. But what's so bad about making the ascribed content salient? Isn't that necessary in order to reject it?

Maxime Lepoutre (2019) connects this discussion to various studies, which appear to show that raising the salience of a piece of ignorant speech, even on the way to rejecting it, itself causes people to be more likely to adopt the false belief.[155] For instance, Stephan Lewandowsky (2012) reports efforts by a health authority to correct myths about the dangers of vaccines by producing leaflets that challenged the myths, but which in fact led those who read them to be more likely to believe them. Such research seems to

---

[153] Horn (1989)

[154] McKeever and Sterken (2021) discuss Haslanger's proposal and suggest some alternative ways of responding instances of stereotyping using generics.

[155] Ignorant and oppressive types of speech are distinct, but the types of verbal expressions of stereotypes using generics that I discuss here are discussed as instances of both.

suggest that making content salient may itself be a problem and hence, that strategies that rely on it are ill-fated.

Where does this leave us with respect to stereotyping generics? At the very least, it should make us nervous about the effects of well-meaning interventions like the ones recommended by Sbisà and Haslanger. If the above effect is a general one, then explicitation, or metalinguistic denial may be not just ineffective, but actively harmful in causing interlocutors to believe stereotypes.

Instead, Lepoutre recommends what he calls 'positive counterspeech' (2019:34). Positive counterspeech is speech that engages with ignorant speech, not by denying it but by making a claim that is incompatible with the ignorant speech, and hence entail its falsity. Applied to generics, the idea would be that instead of responding to utterances of stereotyping by asking for clarification and meta-linguistically denying them, we ought instead to produce speech that itself implicitly conveys that the stereotype is false.

This sounds good, but would, if what I've argued so far is on the right track, be exceedingly difficult in the case of generics. My argument in this thesis has been that generics are non-specific, which in turn makes them easy to make true, believe, and know. In particular, my argument about negated generics in Chapter 3 suggests that it will take all salient generalisation to be false to make a relevant generic false. Even when this occurs, it would be very difficult to convey that fact without repeating various non-generic generalisations that could, in the case of stereotyping generics, cause one's interlocutors to believe them.

To end on a slightly less bleak note, here is a different idea. I've argued that what gives generics their potency is that they are inferentially powerful on the basis of limited evidence. This mechanism relied on the idea that generic beliefs dispose those who hold them to expect that if at least one non-generic about a kind and property is true, then so are several. However, I also argued that background beliefs can disrupt these inference patterns. For example, if we know that less than half of all ducks can give birth to young in any way (fertile female ducks) then we don't assume that, just because *some* generalisation about ducks and laying eggs must be true, the generalisation that *most* ducks lay eggs must be true. The background beliefs we have about kinds shape the types of inferences that our generic beliefs about these kinds dispose us to make.

In the case of social groups, we often lack such background knowledge, and indeed have strong essentialist biases that affirm co-occurrence expectations in their absence. These might be responsible for the especially strong co-occurrence expectations arising from generic beliefs about social kinds, which are reflected in the essentialisation data Leslie draws attention to.

If we could disrupt the mechanisms responsible for strong co-occurrence expectations that generic beliefs about social kids give rise to, then we could dull the potency of stereotypes. Addressing co-occurrence expectations would also be helpful as a way of responding to stereotyping in a way that doesn't just

consist in telling people to just not verbally express the stereotypes they hold, but to lessen their influence.[156]

Strategies for this might take the form of simply increasing the number of accurate beliefs about members of social kinds, and thus make us less likely to allow generic beliefs to fill the gaps and give rise to false and harmful beliefs about them. Alternatively, we might address co-occurrence expectations directly by focussing efforts on ensuring that people understand that just because a property is prevalent in a group, that does not entail that members of the kind share a common nature which *necessitates* that property or makes it right that they have the property. The hope is that making us more generalisation-literate will help us better distinguish co-occurrence expectations that are well-founded from those that are not, and therefore prevent forming false beliefs about social groups.

8 Conclusion

In this last chapter, I provided an account of the role generics play in stereotyping. I argued that stereotypes involve generic beliefs alongside other non-propositional content. Such beliefs are, according to the generality account, by default non-specific. I've argued how this can explain the ease with which speakers can shrug off calls for retraction or use generics to test alignment with others. However, the downside of generic beliefs in the social domain is that they dispose us to faulty and harmful reasoning patterns. As a result, using generics to express stereotypes constitutes a type of misleading speech by negligently conveying harmful falsehoods. I've ended the chapter by considering what my analysis can tell us about how we should respond to stereotyping generics.

[156] Here, I echo Saul (2016) who argues that " […] we shouldn't try to avoid generics. Instead, we need to get better at talking and thinking about them. We need to press people to spell out their evidence for their generic claims and to reflect on what evidence really does or doesn't warrant." (2016: 14)

Conclusion

Looking back

In this thesis, I have presented a novel account of genericity and argued that it is well placed to explain non-specificity in generics, and with this, functional features that depend on it. I developed my arguments over six chapters:

In Chapter 1, I considered how generics should be defined. I evaluated two commonly used ways of introducing generics and argued that they cannot be used to distinguish between generics and overtly quantified generalisations. Instead, I argued that generics should be defined directly as generalisations that do not contain quantifier expressions at the surface level.

In Chapter 2, I discussed the truthmaker variability that this lack of markedness gives rise to. I argued that an underappreciated feature of generics is their non-specificity. I evaluated various ways of accommodating it, including as a by-product of context-sensitivity or semantic incompleteness. However, I argued that, instead, we should consider generics as semantically expressing non-specific generalisations.

In Chapter 3, I presented just such a view. According to the generality account, generics are true if and only if at least one of several non-generic generalisations about the kind and property in question is true. Although generics have weak semantic contents the same cannot be said for their pragmatic effects.

In Chapter 4, I considered Sarah-Jane Leslie's account of generics as giving voice to cognitively primitive generalisations but argued that it was empirically unmotivated and theoretically awkward. As an alternative, I extended the generality account to generic beliefs. The same non-specific generalisations that can be expressed by generic sentences can also be the objects of propositional attitudes, including of belief.

In Chapter 5, I explored a set of interrelated functions of genericity. I argued that because generic beliefs are easily made true, they are also easily believed, easily justified, and easily asserted in comparison to beliefs in overtly quantified generalisations. Generic generalisations thus make it possible to form justified, inferentially powerful beliefs on the basis of limited evidence and communicate them via assertions.

In Chapter 6, I took a closer look at the connection between generics and stereotyping. I gave a novel account of stereotypes as involving generic beliefs alongside non-propositional content. Generics enable speakers to categorise social groups based on limited evidence and communicate such generalisations while also retaining plausible deniability. I ended by considering how this account might inform our responses to verbal stereotyping.

My aim has been to develop a novel account of genericity based on the idea that they are non-specific and demonstrate its utility by considering how non-specificity can illuminate the role of genericity in various phenomena. I hope that the arguments in this thesis are attractive enough to warrant serious consideration of non-specificity approaches to generics and applications of them to other domains. I want to end by sketching what further work I believe would be useful, both in evaluating my proposal and developing it further.

Further work

In Chapter 1, I defined generics by their lack of overt quantifier expression. I suggested that this lack of markedness led to widespread use variability. This claim should be evaluated by considering cross-linguistic data about other languages, specifically ones which do not allow for bare, i.e., non-marked, noun phrase constructions. I also suggested that some types of overtly quantified generalisations may display more use variability than others and that they may therefore be viewed as possessing a degree of genericity. This claim must also be substantiated by taking a closer look at the overtly quantified generalisations that do exhibit some use variability, and the ways in which that variability differs from that displayed by generics.

In Chapter 2, I considered some meta-semantic explanations of non-specificity, including propositional pluralism. My claim that the generality account provides a more straightforward and unified understanding of non-specificity, including in the propositional attitudes towards the propositions expressed by generics, depends for its plausibility on there not being a comparable account that makes use of propositional pluralism. An integration of pluralism with propositional attitudes like belief could make propositional pluralism a much stronger contender.

In Chapter 3, I argued that non-specificity is encoded semantically by means of existential quantification. This predicts that generics have very weak truth-conditions, something that could be tested empirically. Such work could also be of use in developing non-specificity approaches by providing us with a better understanding of *which* non-generic generalisations generics quantify over. I presented the generality account as an account of paradigmatic bare plural generics. Nevertheless, it would be good to consider whether the account could be extended to cover indefinite singulars, definite singular, or even habitual sentences like *Paula gets croissants from the bakery on the weekend*, which have also independently been given accounts that involve existential closure. [157]

In Chapter 4, I extended the generality account to cover generic beliefs, and compared it to Leslie's approach. I advocated for the generality account, partly on the basis that current experimental data does not support Leslie's view conclusively, and partly on the basis that it can provide a unified view of genericity in cognition and language. Should future empirical work provide further support for generic

[157] See Boneh on habituals (2012) and Corblin on IS generics (2012).

beliefs as a distinctive System 1 generalisations, then the account I proposed could be integrated with Leslie's view of generics as cognitively primitive, while maintaining that the content of generic beliefs is also that of generic sentences, hence avoiding the disconnect between genericity in mental states and as a feature of sentences in Leslie's account. I also argued that generic beliefs are inferentially powerful by way of disposing those who hold them to expect that different non-generic generalisations co-occur. This hypothesis could also be tested empirically and linked to the vast literature on different kinds of cognitive biases that might explain those expectations.

In Chapter 5, I gave an account of an epistemic function of genericity. I suggested that my argument might explain how the heuristics through which we apply concepts, or the application conditions of the concepts themselves, were generic in form. This line of argument could be developed through closer engagement with the literature on concepts and concept formation as well as on what has been called 'definitional' uses of generics.

In Chapter 6, I applied the analysis of generic beliefs from previous chapters to the phenomenon of stereotyping. I argued that even though stereotypes expressed using generics may sometimes be true, they nevertheless dispose those who hold them to false, or otherwise harmful beliefs. I suggested that this may be partly because of especially strong co-occurrence expectations for social kinds, something we may be able to disrupt. This proposal would benefit from closer engagement with the literature on how we represent beliefs about social kinds to gain a better understanding of if and how co-occurrence biases come about with respect to generic beliefs about social groups.

Bibliography

Adebayo, Taofeeq. 2018. "Genericity in Event Semantics: A Look at Yoruba Generic Sentences." In *Proceedings of Triple A 4: Fieldwork Perspectives on the Semantics of African, Asian and Austronesian Languages*, edited by Bogal-Allbritten Elizabeth Coppock Elizabeth, 48–64.

Anderson, Luvell, Sally Haslanger, and Rae Langton. 2012. "Language and Race." In *Routledge Companion to Philosophy of Language*. Routledge.

Antinucci, Francesco, and Ruth Miller. 1976. "How Children Talk about What Happened." *Journal of Child Language* 3 (2): 167–89.

Asher, Nicholas, and Michael Morreau. 1995. "What Some Generic Sentences Mean." In *The Generic Book*, edited by Gregory Carlson and Francis Jeffry Pelletier, 300–338. University of Chicago Press Chicago.

Asher, Nicholas, and Francis Jeffry Pelletier. 2012. "More Truths about Generic Truth." In *Genericity*, 313–33. Oxford University Press.

Assarsson, Elias. 2012. "A Two-Tiered Theory of Generics: A Synthesis of Leslie's and Liebesman's Theories." Lund University.

Bach, Kent. 1994. "Conversational Impliciture." *Mind & Language* 9 (2): 124–62.

———. 2001a. "You Don't Say?" *Synthese* 128 (1): 15–44.

———. 2001b. "Speaking Loosely: Sentence Nonliterality." *Midwest Studies in Philosophy* 25 (1): 249–63.

———. 2006. "The Excluded Middle: Semantic Minimalism without Minimal Propositions." *Philosophy and Phenomenological Research* 73 (2): 435–42.

———. 2005. "The Emperor's New 'Knows.'" In *Contextualism in Philosophy*, edited by Gerhard Preyer and Georg Peter. Clarendon Press. Accessed July 7, 2021.

Barwise, Jon, and Robin Cooper. 1981. "Generalized Quantifiers and Natural Language." *Linguistics and Philosophy* 4 (2): 159–219.

Beaver, David I., Craige Roberts, Mandy Simons, and Judith Tonhauser. 2017. "Questions Under Discussion: Where Information Structure Meets Projective Content." *Annual Review of Linguistics* 3 (1): 265–84.

Beebe, James. 2012. "Social Functions of Knowledge Attributions." In *Knowledge Ascriptions*, edited by Mikkel Gerken and Jessica Brown, 220–42. Oxford University Press.

Beeghly, Erin. 2015. "What Is a Stereotype? What Is Stereotyping?" *Hypatia* 30 (4): 675–91.

Benton, Matthew A., and Peter Van Elswyk. n.d. "Hedged Assertion." In *The Oxford Handbook of Assertion*, edited by Sanford Goldberg, 245–63. Oxford University Press.

Bhatt, Rajesh. 2006. *Covert Modality in Non-Finite Contexts*. Walter de Gruyter.

Blum, Lawrence. 2004. "Stereotypes And Stereotyping: A Moral Analysis." *Philosophical Papers* 33 (3): 251–89.

Boneh, Nora. 2012. "Hab and Gen in the Expression of Habituality." In *Genericity*, edited by Alda Mari, Claire Beyssade, and Fabio Del Prete, 176–91. Oxford University Press.

Brandone, Amanda C., Andrei Cimpian, Sarah-Jane Leslie, and Susan A. Gelman. 2012. "Do Lions Have Manes? For Children, Generics Are about Kinds rather than Quantities." *Child Development* 83 (2): 423–33.

Burton-Roberts, Noel. 1976. "On the Generic Indefinite Article." *Language* 52 (2): 427–48.

Cappelen, Herman. 2018. "On Appeals to a Concept's Purpose or Function." In *Fixing Language: An Essay on Conceptual Engineering*. Oxford University Press.

Carlson, Greg. 1977. "Reference to Kinds in English." University of Massachusetts, Amherst.

———. 1989. "On the Semantic Composition of English Generic Sentences." In *Properties, Types and Meaning: Volume II: Semantic Issues*, edited by Gennaro Chierchia, Barbara H. Partee, and Raymond Turner, 167–92. Dordrecht: Springer Netherlands.

———. 1995. "Truth Conditions of Generic Sentences: Two Contrasting Views." *The Generic Book*, 224–37.

———. 2008. "Patterns in the Semantics of Generic Sentences." In *Time and Modality*, edited by Jacqueline Guéron and Jacqueline Lecarme, 17–38. Dordrecht: Springer Netherlands.

Chignell, Andrew. 2018. "The Ethics of Belief." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. https://plato.stanford.edu/archives/spr2018/entries/ethics-belief/.

Cimpian, Andrei, Amanda C. Brandone, and Susan A. Gelman. 2010. "Generic Statements Require Little Evidence for Acceptance but Have Powerful Implications." *Cognitive Science* 34 (8): 1452–82.

Cimpian, Andrei, and Lucy C. Erickson. 2012. "Remembering Kinds: New Evidence That Categories Are Privileged in Children's Thinking." *Cognitive Psychology* 64 (3): 161–85.

Cimpian, Andrei, and Ellen M. Markman. 2008. "Preschool Children's Use of Cues to Generic Meaning." *Cognition* 107 (1): 19–53.

Cimpian, Andrei, Trent J. Meltzer, and Ellen M. Markman. 2011. "Preschoolers' Use of Morphosyntactic Cues to Identify Generic Sentences: Indefinite Singular Noun Phrases, Tense, and Aspect." *Child Development* 82 (5): 1561–78.

Cohen, Ariel. 1999a. "Think Generic! The Meaning and Use of Generic Sentences." Carnegie Mellon University.

———. 1999b. "Generics, Frequency Adverbs, and Probability." *Linguistics and Philosophy* 22 (3): 221–53.

———. 2001a. "Relative Readings of Many, Often, and Generics." *Natural Language Semantics* 9 (1): 41–67.

———. 2001b. "On the Generic Use of Indefinite Singulars." *Journal of Semantics* 18 (3): 183–209.

———. 2004. "Generics and Mental Representations." *Linguistics and Philosophy* 27 (5): 529–56.

———. 2012. "No Quantification without Reinterpretation." In *Genericity*, edited by Alda Mari, Claire Beyssade, and Fabio Del Prete. Oxford University Press.

Collins, John. 2018. "Genericity sans Gen." *Mind & Language* 33 (February): 34–64.

Corblin, Francis. 2012. "The Roots of Genericity: Indefinite Singulars vs Definite Plurals∗." In *Genericity*, 352–71. Oxford University Press.

Craig, Edward. 1999. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Clarendon Press.

Cromer, Richard F. 1968. "The Development of Temporal Reference During the Acquisition of Language." Harvard University.

Dahl, Östen. 1975. "On Generics." In *Formal Semantics of Natural Language*, edited by Edward L. Keenan, 99–111. Cambridge University Press.

———. 1995. "The Marking of the Episodic/generic Distinction in Tense-Aspect Systems." In *The Generic Book*, edited by Greg Carlson and Francis Jeffry Pelletier, 412–25. Chicago University Press.

Delgrande, James P. 1987. "A First-Order Conditional Logic for Prototypical Properties." *Artificial Intelligence* 33 (1): 105–30.

Diesing, Molly. 1992. *Indefinites*. MIT press.

Eckardt, Regine. 1999. "Normal Objects, Normal Worlds and the Meaning of Generic Sentences." *Journal of Semantics* 16 (3): 237–78.

Fara, Delia Graff. 2015. "Descriptions as Predicates." *The Philosophical Review* 124 (1): 59–117.

Fintel, Kai von. 1994. "Restrictions on Quantifier Domains." University of Massachusetts at Amherst.

Fintel, Kai von, and Anthony Gillies. 2011. "'Might' Made Right." In *Epistemic Modality*, edited by Andy Egan and Brian Weatherson. Oxford Scholarship Online.

Fodor, Jerry A. 1975. *The Language of Thought*. Harvard University Press.

Frankish, Keith. 2010. "Dual-Process and Dual-System Theories of Reasoning: Dual-Process and Dual-System Theories of Reasoning." *Philosophy Compass* 5 (10): 914–26.

Frederick, Shane. 2005. "Cognitive Reflection and Decision Making." *The Journal of Economic Perspectives: A Journal of the American Economic Association* 19 (4): 25–42.

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

Gelman, Susan A., Peggy J. Goetz, Barbara W. Sarnecka, and Jonathan Flukes. 2008. "Generic Language in Parent-Child Conversations." *Language Learning and Development: The Official Journal of the Society for Language Development* 4 (1): 1–31.

Gelman, Susan A., and Lakshmi Raman. 2003. "Preschool Children Use Linguistic Form Class and Pragmatic Cues to Interpret Generics." *Child Development* 74 (1): 308–25.

Gelman, Susan A., Elizabeth A. Ware, and Felicia Kleinberg. 2010. "Effects of Generic Language on Category Content and Structure." *Cognitive Psychology* 61 (3): 273–301.

Gibbard, Allan. 2009. *Thinking How to Live*. Harvard University Press.

Goldin-Meadow, Susan, Susan A. Gelman, and Carolyn Mylander. 2005. "Expressing Generic Concepts with and without a Language Model." *Cognition* 96 (2): 109–26.

Greenberg, Yael. 2004. *Manifestations of Genericity*. Taylor & Francis.

———. 2007. "Exceptions to Generics: Where Vagueness, Context Dependence and Modality Interact." *Journal of Semantics* 24 (2): 131–67.

Hammond, Matthew D., and Andrei Cimpian. 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): 607–14.

Haslam, Nick, Louis Rothschild, and Donald Ernst. 2002. "Are Essentialist Beliefs Associated with Prejudice?" *The British Journal of Social Psychology / the British Psychological Society* 41 (Pt 1): 87–100.

Haslanger, Sally. 2011. "Ideology, Generics, and Common Ground." In *Feminist Metaphysics: Explorations in the Ontology of Sex, Gender and the Self*, edited by Charlotte Witt, 179–207. Dordrecht: Springer Netherlands.

———. 2014. "The Normal, the Natural and the Good: Generics and Ideology" 365 (September): 392.

Haspelmath, Martin. 2006. "Against Markedness (And What to Replace It With)." *Journal of Linguistics* 42 (1): 25–70.

Heim, Irene. 1982. "The Semantics of Definite and Indefinite Noun Phrases." UMass Amherst.

Hoicka, Elena, Jennifer Saul, Eloise Prouten, Laura Whitehead, and Rachel Sterken. 2018. "Language Signalling High Proportions, Not Just Generics, Leads to Essentializing for Novel Social Kinds." *Psyarxiv.com*.

Hollander, Michelle A., Susan A. Gelman, and Jon Star. 2002. "Children's Interpretation of Generic Noun Phrases." *Developmental Psychology* 38 (6): 883–94.

Horn, Laurence. 1989. *A Natural History of Negation*. University of Chicago Press.

———. 2020. "Neg-Raising." In *The Oxford Handbook of Negation*, edited by Viviane Déprez And Teresa. Oxford University Press.

Jackson, Frank, and Philip Pettit. 1998. "A Problem for Expressivism." *Analysis* 58 (4): 239–51.

Johnston, Mark, and Sarah-Jane Leslie. 2012. "Concepts, Analysis, Generics and the Canberra Plan." *Philosophical Perspectives. A Supplement to Nous* 26 (1): 113–71.

Jussim, Lee, Jarret T. Crawford, and Rachel S. Rubinstein. 2015. "Stereotype (In)Accuracy in Perceptions of Groups and Individuals:" *Current Directions in Psychological Science*, December.

Kahneman, Daniel. 2002. "Mapping Bounded Rationality: A Perspective on Intuitive Judgment and Choice." Presented at the Nobel Lecture, Stockholm, Sweden, December 8.

———. 2011. *Thinking, Fast and Slow*. Macmillan.

Kahneman, Daniel, and Shane Frederick. 2002. "Representativeness Revisited: Attribute Substitution in Intuitive Judgment." In *Heuristics and Biases: The Psychology of Intuitive Judgment*, edited by Thomas Gilovich, Dale Griffin, and Daniel Kahneman, 49–81. Cambridge University Press.

Kamp, Hans. 2002 (previously published in 1981). "A Theory of Truth and Semantic Representation." In *Formal Semantics: The Essential Readings*, edited by Paul Portners and Barbara H. Partee, 277–322. Blackwell.

Khemlani, Sangeet, Sarah-Jane Leslie, and Sam Glucksberg. 2009. "Generics, Prevalence, and Default Inferences." *Proceedings of the Cognitive Science Society*, 443–48.

———. 2012. "Inferences about Members of Kinds: The Generics Hypothesis." *Language and Cognitive Processes* 27 (6): 887–900.

King, Jeffrey C. 2013. "Supplementives, the Coordination Account, and Conflicting Intentions." *Philosophical Perspectives. A Supplement to Nous* 27 (1): 288–311.

Kirkpatrick, James Ravi. 2019. "Essays on Genericity." University of Oxford.

Kratzer, Angelika. 1977. "What 'Must' and 'Can' Must and Can Mean." *Linguistics and Philosophy* 1 (3): 337–55.

———. 2012. *Modals and Conditionals: New and Revised Perspectives*. OUP Oxford.

Kratzer, Angelika, Hans-Jürgen Eikmeyer, and Hannes Rieser. 1981. "The Notional Category of Modality." *Formal Semantics: The Essential Readings*, 289–323.

Krifka, Manfred. 2003. "Bare NPs: Kind-Referring, Indefinites, Both, or Neither?" *Semantics and Linguistic Theory* 13 (0): 180–203.

Krifka, Manfred, Francis Jeffry Pelletier, Gregory Carlson, Alice ter Meulen, Gennaro Chierchia, and Godehard Link. 1995. "Genericity: An Introduction." In *The Generic Book*, edited by Greg Carlson and Francis Jeffry Pelletier, 1–124. University of Chicago Press.

Kusch, Martin. 2009. "Testimony and the Value of Knowledge." In *Epistemic Value*, edited by Pritchard, Haddock, and MIllar, 60–94. Oxford: Oxford University Press.

Kusch, Martin, and Robin McKenna. 2020. "The Genealogical Method in Epistemology." *Synthese* 197 (3): 1057–76.

Kvanvig, Jonathan. 2009. "Assertion, Knowledge, and Lotteries." In *Williamson on Knowledge*, edited by Duncan Pritchard and Patrick Greenough, 140–60. Oxford: Oxford University Press.

Lackey, Jennifer. 2006. "Knowing From Testimony." *Philosophy Compass* 1 (5): 432–48.

———. 2007. "Norms of Assertion." *Noûs* 41 (4): 594–626.

Langton, Rae. 2018. "Blocking as Counter-Speech." In *New Work on Speech Acts*, edited by Daniel Fogal, Daniel W. Harris, and Matt Moss. Oxford University Press.

Lawler, John Michael. 1973. "Studies in English Generics." University of Michigan.

Lazaridou-Chatzigoga, Dimitra. 2019. "Genericity." In *The Oxford Handbook of Experimental Semantics and Pragmatics*, edited by Chris Cummins and Napoleon Katsos. Oxford University Press.

Lazaridou-Chatzigoga, Dimitra, Napoleon Katsos, and Linnaea Stockall. 2015. "Genericity Is Easy? Formal and Experimental Perspectives." *Ratio* 28 (4): 470–94.

Lazaridou-Chatzigoga, Dimitra, and Linnaea Stockall. 2013. "Genericity, Exceptions and Domain Restriction: Experimental Evidence from Comparison with Universals." *Proceedings of Sinn Und Bedeutung* 17: 325–43.

Lazaridou-Chatzigoga, Dimitra, Linnaea Stockall, and Napoleon Katsos. 2017. "A New Look at the 'Generic Overgeneralisation' Effect." *Inquiry: A Journal of Medical Care Organization, Provision and Financing* 0 (0): 1–27.

———. 2019. "Contextualising Generic and Universal Generalisations: Quantifier Domain Restriction and the Generic Overgeneralisation Effect." *Journal of Semantics* 36 (4): 617–64.

Lemeire, Olivier. 2020. "Falsifying Generic Stereotypes." *Philosophical Studies* 178 (7): 2293–2312.

Lepoutre, Maxime. 2019. "Can 'More Speech' Counter Ignorant Speech?" *Journal of Ethics and Social Philosophy* 16 (3).

Lerner, Adam, and Sarah-Jane Leslie. 2013. "Generics, Generalism, and Reflective Equilibrium: Implications for Moral Theorizing from the Study of Language." *Philosophical Perspectives. A Supplement to Nous* 27 (1): 366–403.

Leslie, Sarah-Jane. 2007. "Generics and the Structure of the Mind." *Philosophical Perspectives. A Supplement to Nous* 21 (1): 375–403.

———. 2008. "Generics: Cognition and Acquisition." *The Philosophical Review* 117 (1): 1–47.

———. 2012. "Generics articulate default generalizations." *Recherches linguistiques de Vincennes*, no. 41: 25–44.

———. 2013. "Essence and Natural Kinds: When Science Meets Preschooler Intuition 1." In *Oxford Studies in Epistemology Volume 4*. Oxford: Oxford University Press.

———. 2014. "Carving Up the Social World with Generics*." In *Oxford Studies in Experimental Philosophy*, 208–31. Oxford: Oxford University Press.

———. 2015. "'Hillary Clinton Is the Only Man in the Obama Administration': Dual Character Concepts, Generics, and Gender." *Analytic Philosophy* 56 (2): 111–41.

———. 2017. "The Original Sin of Cognition: Fear, Prejudice, and Generalization." *The Journal of Philosophy* 114 (8): 393–421.

Leslie, Sarah-Jane, and Susan A. Gelman. 2012. "Quantified Statements Are Recalled as Generics: Evidence from Preschool Children and Adults." *Cognitive Psychology* 64 (3): 186–214.

Leslie, Sarah-Jane, Sangeet Khemlani, and Sam Glucksberg. 2011. "Do All Ducks Lay Eggs? The Generic Overgeneralization Effect." *Journal of Memory and Language* 65 (1): 15–31.

Leslie, Sarah-Jane, and Adam Lerner. 2016. "Generic Generalizations." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2016. y.

Lewis, David K. 1972. "Psychophysical and Theoretical Identifications." *Australasian Journal of Philosophy* 50 (3): 249–58.

———. 1980. "Mad Pain and Martian Pain." In *Readings in the Philosophy of Psychology*, edited by Ned Block, 216–22. Harvard University Press.

———. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74 (4): 549–67.

Liebesman, David. 2011. "Simple Generics." *Noûs* 45 (3): 409–42.

Liebesman, David, and Ofra Magidor. 2017. "Copredication and Property Inheritance." *Philosophical Issues. A Supplement to Nous* 27 (1): 131–66.

Link, Godehard. 1987. "Generalized Quantifiers and Plurals." In *Generalized Quantifiers*, 151–80. Dordrecht: Springer Netherlands.

MacFarlane, John. 2005. "The Assessment Sensitivity of Knowledge Attributions." In *Oxford Studies in Epistemology*, edited by Tamar Szabó Gendler and John Hawthorne, 197–234. Oxford: Oxford University Press.

Maitra, Ishani, and Brian Weatherson. 2010. "Assertion, Knowledge, and Action." *Philosophical Studies* 149 (1): 99–118.

Mari, Alda, Claire Beyssade, and Fabio Del Prete. 2012. "Introduction." In *Genericity*, edited by Alda Mari, Claire Beyssade, and Fabio Del Prete. Oxford Studies in Theoretical Linguistics. Oxford: Oxford

University Press.

Mcconnell-Ginet, Sally. 2012. "Generic Predicates and Interest-Relativity." *Canadian Journal of Linguistics/Revue Canadienne de Linguistique* 57 (2): 261–87.

McCook, Alison. 2017. "'I Placed Too Much Faith in Underpowered Studies:' Nobel Prize Winner Admits Mistakes." February 20, 2017. https://retractionwatch.com/2017/02/20/placed-much-faith-underpowered-studies-nobel-prize-winner-admits-mistakes/.

McCready, Elin. 2014. *Reliability in Pragmatics*. Oxford University Press.

McGrath, Matthew. 2015. "Two Purposes of Knowledge Attribution and the Contextualism Debate." In *Epistemic Evaluation*, edited by John Greco and David Henderson. Oxford University Press.

McKeever, Matthew, and Rachel Sterken. 2021. "Social and Political Aspects of Generic Language and Speech." In *The Routledge Handbook of Social and Political Philosophy of Language*, 259–80. Routledge.

Meyer, Meredith, Susan A. Gelman, and Sarah M. Stilwell. 2011. "Generics Are a Cognitive Default: Evidence from Sentence Processing." In *Proceedings of the Cognitive Science Society*. Vol. 33.

Milsark, Gary. 1974. "Existential Sentences in English." Massachusetts Institute of Technology.

Munton, Jessie. 2019. "Beyond Accuracy: Epistemic Flaws with Statistical Generalizations." *Philosophical Issues. A Supplement to Nous* 29 (1): 228–40.

Nguyen, Anthony. 2019. "The Radical Account of Bare Plural Generics." *Philosophical Studies* 177 (5): 1303–31.

Nickel, Bernhard. 2006. "On Semantics for Characterizing Sentences." In *Recent Work on Modality*, edited by J. Gajewski, V. Hacquard, B. Nickel, & S. Yalcin, 123–47. MIT Press.

———. 2008. "Generics and the Ways of Normality." *Linguistics and Philosophy* 31 (6): 629–48.

———. 2010. "Ceteris Paribus Laws: Generics and Natural Kinds." *Philosopher's Imprint* 10 (06).

———. 2016. *Between Logic and the World: An Integrated Theory of Generics*. Oxford University Press UK.

———. 2018. "Generics, Conservativity, and Kind-Subordination." *Philosophers' Imprint* 18 (13): 18.

O'Donnell, Patrick. 2017. "Generics, Race, and Social Perspectives." *Inquiry: An Interdisciplinary Journal of Philosophy*, 1–38.

Osman, Magda. 2004. "An Evaluation of Dual-Process Theories of Reasoning." *Psychonomic Bulletin & Review* 11 (6): 988–1010.

Pelletier, Francis Jeffry. 2009. "Generics: A Philosophical Introduction." In *Kinds, Things, and Stuff*, 3–15. Oxford University Press.

Pelletier, Francis Jeffry, and Nicholas Asher. 1997. "Generics and Defaults." In *Handbook of Logic and*

Language, edited by Johan van Benthem and Alice G. B. ter Meulen, 1125–77. Elsevier.

Prasada, Sandeep. 2000. "Acquiring Generic Knowledge." *Trends in Cognitive Sciences* 4 (2): 66–72.

Prasada, Sandeep, and Elaine M. Dillingham. 2006. "Principled and Statistical Connections in Common Sense Conception." *Cognition* 99 (1): 73–112.

Prasada, Sandeep, Sangeet Khemlani, Sarah-Jane Leslie, and Sam Glucksberg. 2013. "Conceptual Distinctions amongst Generics." *Cognition* 126 (3): 405–22.

Preston, Beth. 1998. "Why Is a Wing Like a Spoon? A Pluralist Theory of Function." *The Journal of Philosophy* 95 (5): 215–54.

Putnam, Hilary. 1963. "Brains and Behavior." In *Analytical Philosophy: Second Series*, edited by Ronald J. Butler. Blackwell.

Queloz, Matthieu. 2021. *The Practical Origins of Ideas: Genealogy as Conceptual Reverse-Engineering (Open Access)*. Oxford: Oxford University Press.

Rhodes, Marjorie, Sarah-Jane Leslie, and Christina M. Tworek. 2012. "Cultural Transmission of Social Essentialism." *Proceedings of the National Academy of Sciences of the United States of America* 109 (34): 13526–31.

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

Roberts, Craige. 1989. "Modal Subordination and Pronominal Anaphora in Discourse." *Linguistics and Philosophy* 12 (6): 683–721.

———. 2012. "Information Structure: Towards an Integrated Formal Theory of Pragmatics." *Semantics and Pragmatics* 5 (0): 6–1 – 69.

Rosola, Martina. 2019. "Generics in Context: The Robustness and the Explanatory Implicatures." In *Modeling and Using Context*, edited by Gábor Bella and Paolo Bouquet, 11939:223–37. Cham: Springer International Publishing.
———. 2021. "The Harmful Feature of Generics." Università degli Studi di Genova.

Rysiew, Patrick. 2012. "Epistemic Scorekeeping." In *Knowledge Ascriptions*, edited by Jessica Brown and Mikkel Gerken. Oxford University Press.

Saul, Jennifer. 2017. "Are Generics Especially Pernicious?" *Inquiry*, 1–18.

———. 2018. "Negligent Falsehood, White Ignorance, and Fake News." In *Lying: Language, Knowledge, Ethics, and Politics*. Engaging Philosophy. Oxford, New York: Oxford University Press.

Sbisà, Marina. 1999. "Ideology and the Persuasive Use of Presupposition." In *Language and Ideology. Selected Papers from the 6th International Pragmatics Conference*, edited by Jef Verschueren, 492–509. 1. Antwerp: Internatonal Pragmatics Association.

Shoemaker, Sydney. 1981. "Some Varieties of Functionalism." *Philosophical Topics* 12 (1): 93–119.

Smith, Eliot R., and Jamie DeCoster. 2000. "Dual-Process Models in Social and Cognitive Psychology: Conceptual Integration and Links to Underlying Memory Systems." *Personality and Social Psychology Review: An Official Journal of the Society for Personality and Social Psychology, Inc* 4 (2): 108–31.

Spector, Benjamin. 2016. "Comparing Exhaustivity Operators." *Semantics & Pragmatics* 9 (11).

Stalnaker, Robert. 1978. "Assertion." *Syntax and Semantics (New York Academic Press)* 9: 315–32.

Stanley, Jason, and Zoltan Gendler Szabo. 2000. "On Quantifier Domain Restriction." *Mind & Language* 15 (2&3): 219–61.

Sterken, Rachel Katharine. 2015a. "Generics in Context." *Philosophers' Imprint* 15 (21): 30.

———. 2015b. "Leslie on Generics." *Philosophical Studies* 172 (9): 2493–2512.

———. 2016. "Generics, Covert Structure and Logical Form." *Mind & Language* 31 (5): 503–29.

Teichman, Matt. 2016. "The Sophisticated Kind Theory." *Inquiry: A Journal of Medical Care Organization, Provision and Financing* 0 (0): 1–47.

Thomasson, Amie L. 2020. "A Pragmatic Method for Normative Conceptual Work." In *Conceptual Engineering and Conceptual Ethics*, edited by Alexis Burgess, Herman Cappelen, and David Plunkett. Oxford University Press.

Viebahn, Emanuel. 2019. *Semantic Pluralism*. Klostermann.

Visser, Alnica. n.d. "Genericity Generalized." University of Pittsburgh.

Wasserman, Ryan. 2011. "Dispositions and Generics." *Philosophical Perspectives. A Supplement to Nous* 25 (1): 425–53.

Whiteley, Ella. (forthcoming) "Harmful Salience Perspectives." In *Salience: A Philosophical Inquiry*, edited by Sophie Archer. Routledge.

Williamson, Timothy. 1996. "Knowing and Asserting." *The Philosophical Review* 105 (4): 489.

———. 2005. "Knowledge, Context, and the Agent's Point of View." In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, edited by Gerhard Preyer and Georg Peter, 91–114. Oxford University Press.

Wittgenstein, Ludwig. 1922. *Tractatus Logico-Philosophicus*. Kegan Paul, Trench, Trubner & Co. London, New York: Harcourt

Wodak, Daniel, Sarah-Jane Leslie, and Marjorie Rhodes. 2015. "What a Loaded Generalization: Generics and Social Cognition." *Philosophy Compass* 10 (9): 625–35. *Compass* 10 (9): 625–35.