"Philosophers care about the truth": Descriptive/Normative Generics without Ambiguity

Olivier Lemeire – This is a draft, Please do not cite without permission

Abstract

Some generic generalizations have both a descriptive and a normative reading. The generic sentence "Philosophers care about the truth," for instance, can be read as describing what philosophers in fact care about, but can also be read as prescribing philosophers to care about the truth. Similarly, the generic stereotype "Men are tough" can also be understood in a normative way, saying that men ought to be tough. According to one prominent view, these descriptive/normative generics are *ambiguous* between expressing a descriptive generalization and a normative generalization. This paper first argues against this ambiguity thesis, focusing on Leslie's most recent articulation of it. In response, a new contextualist semantic theory for generics is introduced. On this theory, generics express context-dependent kindhood generalizations. Generics for which both a descriptive and a normative reading are available, are contextually underspecified rather than ambiguous.

Introduction

Generic sentences, or *generics* for short, are generalizations that are formulated without the overt use of a quantifier. Typical examples of generics are sentences like the following:

(1) Tigers are striped.

(2) Ravens are black.[1]

On the most widely accepted analysis, generic generalizations like these contain an unpronounced sentential operator in their logical form, called *Gen*. This variable-binding operator relates a restrictor – such as "Tiger(x)" – with a scope – such as "Striped(x)." The logical form of generics like (1) and (2), for example, is generally thought to be this:

(3) Gen x [Tiger(x)][Striped(x)]

(4) Gen x [Raven(x)][Black(x)]

Even after decades of research, there is little agreement in the literature concerning the type of generalization(s) expressed by *Gen* (Krifka et al., 1995; Sterken, 2017). Generics appear to express broad kind-wide generalizations but can still be true in the face of exceptions. The existence of some

[1] Since there appear to be subtle differences in meaning between generics with a bare plural noun phrase and those with an indefinite singular or definite singular noun phrase, I will only discuss the former in this paper (but see: Greenberg, 2002; Krifka et al., 1995).

albino tigers does not falsify (1), for instance, just as the existence of albino ravens does not falsify (2).

Yet their tolerance for exceptions is not the only puzzling fact about the meaning of generics. Another puzzling fact is that in a default context, some generics only have a descriptive reading whereas other generics also have a normative reading. Sentences such as (1) and (2), for instance, are standardly interpreted as just *describing* what members of the kind are like. Other generics, however, can also be interpreted as *prescribing* what members of the kind ought to be like. Here are two examples that will be returned to throughout this paper:

(5) Philosophers care about the truth.

(6) Men are tough.

Generic sentences like (5) and (6) have both a descriptive and a normative reading readily available. Sentence (5), for instance, can be read as describing what philosophers in fact care about, but can also be read as prescribing philosophers to care about the truth. The same holds for sentence (6), which, according to one of its readings, says that men ought to be tough.[2] The question to be answered in this paper is why there is both a descriptive and a normative reading of these sentences so readily available, even in a default context.

More specifically, descriptive/normative generics like (5) and (6) raise the question as to why a normative reading is more readily available for these generics than for generics like (1) and (2). In a default context, there is no reading available of (1) according to which this sentence says that ravens ought to be black in a prescriptive-normative sense. Part of the explanation presumably lies in the different conceptualization of natural kinds such as *ravens* and *tigers* compared to that of human kinds such as *philosophers* and *men.* Unlike human kinds, natural kinds are typically not conceived of as the kind of thing to which prescriptive norms apply. But how exactly does this difference in conceptualization lead to a difference in the readings that are available for generic sentences?

A second relevant question is why generics are particularly prone to having a normative reading. After all, explicitly quantified generalizations about the same kinds that are the subjects of (5) and (6) do not standardly have a normative reading in addition to a descriptive one. Consider, for example, the following two quantified generalizations about philosophers and men:

(7) Most philosophers care about the truth.

(8) All men are tough.

[2] This generic will be mentioned often in this paper. Obviously, this should not be considered an endorsement of it. Other examples of objectionable descriptive/normative generics are “Women are submissive,” “Boys don’t cry,” “Girls are caring,” and “Black people don’t snitch” (from the ABC sitcom *Black-Ish*).

Without further conversational scene-setting, one would not understand (7) as saying that most philosophers *ought* to care about the truth.[3] Instead, it simply expresses a generalization that describes what most philosophers care about. Similarly, without the pragmatic context leading to an alternative interpretation, (8) does not say that all men *ought* to be tough. This sentence has only one reading, stating that all men are in fact tough. An adequate semantic theory for generics should also explain why generics about these same kinds and properties *do* have a normative reading, even in a default context.

In this paper, I argue that generics are prone to having different readings because they express context-dependent kindhood generalizations. Generics, that is, express generalizations that contain *kindhood* as a constituent. I propose that a generic of the form "*Ks* are *F*" says (very roughly) that *F is partly constitutive of the kindhood of Ks*. In addition to this semantic theory for generics, I subscribe to a relational view on kindhood. On this view, *kindhood* can be realized in different ways because it is a property of categories that only exists in relation to a categorizing strategy, and there are many such strategies. As a result of this, a generic sentence is only truth-evaluable given a contextually selected categorizing strategy. Some generic sentences – like (5) and (6) – standardly have two different readings because it is common ground that the category's kindhood can be understood in relation to two different categorizing strategies.

This contextualist theory is defended in response to the dominant view that descriptive/normative generics are *ambiguous* between expressing a descriptive generalization and a normative generalization (Cohen, 2001; Greenberg, 2002; Leslie, 2015). In the first part of the paper, I present two different versions of this ambiguity thesis and an objection against each. Particular attention is paid to Leslie's (2015) recent and influential version of the ambiguity thesis (Del Pinal & Reuter, 2017; Reuter, 2019).[4] In the second part of this paper, I defend my alternative contextualist view on normative generics in response.

[3] This is not to say that quantified sentences cannot convey normative messages in some conversational contexts. If below a mandatory action sign it reads "All dogs wear a collar," this universal generalization conveys a prescription based on pragmatic grounds. Here I simply note that many generics have a normative reading even in a default context whereas similar quantified generalizations do not.

[4] Since submitting this paper, another interesting criticism of Leslie's view has been published, but one pursuing a very different critique and a pragmatic alternative (Hesni, 201). Unfortunately, I cannot discuss this alternative here (but see my SUPRESSED).

1. The Ambiguity Thesis

1.1 Two Ways of Being Ambiguous

The fact that some generic sentences standardly have both a descriptive and a normative reading calls for an explanation. On what is currently the most prominent view, these descriptive/normative generics are *ambiguous* between expressing a descriptive and a normative generalization. In this section, I introduce two different versions of this ambiguity thesis.

According to a first version of the ambiguity thesis, all bare plural generics of the form "*Ks* are *F*" are ambiguous between expressing a descriptive-statistical generalization and an alternative generalization, one that has normative force depending on the kind and property involved (Cohen 2001; Greenberg 2002).[5] As Cohen argues, for instance, all bare plural generics are syntactically ambiguous between a descriptive-statistical generalization and an alternative rule-based generalization. Generics like (1) and (5), for example, are claimed to be ambiguous in the following way:[6]

(9) a. Gen x [Tiger(x)][Striped(x)]

b. In-effect(RULE(Tiger(x) → Striped(x)))

(10) a. Gen x [Philosopher(x)][Cares-about-the-truth(x)]

b. In-effect(RULE(Philosopher(x) → Cares-about-the-truth(x)))

On their a-readings, generics have the logical form standardly attributed to generics, namely one containing an unpronounced quantificational operator called *Gen*. Cohen proposes that *Gen* has probabilistic meaning so that (9a), for example, says that any particular tiger is likely to be striped. Similarly, (10a) says that any particular philosopher probably cares about the truth.

Yet Cohen further argues that bare plural generics can also be used to express a very different type of proposition, one on which they predicate that a rule is in effect according to which members of the kind instantiate the generalized property. Whether this b-reading has normative force depends on the type of rule the speaker is referring to. On its rule-based (b-)reading, "Ravens are

---

[5] This first version of the ambiguity thesis is defended by Cohen (2001) and Greenberg (2002) in the first place to explain the subtle difference in meaning between generics with a bare plural (BP) noun phrase (e.g., Tigers are striped) and generics with an indefinite singular (IS) noun phrase (e.g., A tiger is striped). Whereas BP generics are said to be ambiguous, IS generics are not; they cannot express a descriptive-statistical generalization. By rejecting the ambiguity of BP generics, as I argue we should, one also incurs the burden of having to provide an alternative explanation for the difference in meaning between BP and IS generics. Since this topic falls outside of the scope of this paper, I cannot meet this burden here.

[6] Cohen himself represents the logical form of the rule-based reading as "In-effect(!(Tiger(x) → Striped(x)))" where ! is an operator that maps a formula of the form (K(x)→F(x)) to the rule that it describes. I refer to this !-operator with 'RULE' only to improve the readability of the main text.

black" does not have normative force because the rule that is referred to is standardly taken as some sort of biological law, based on common knowledge about the nature of this biological kind. A generic like "Philosophers care about the truth," however, does have normative force on its rule-based (b-)reading, because one standardly understands the rule that is being referred to as a social one with normative force.

On the first version of the ambiguity thesis then, all generics of the form "*Ks* are *F*" are ambiguous between expressing a descriptive-statistical generalization and a rule-based predication, the latter of which can have normative force depending on the type of rule the speaker is referring to (Cohen 2001). Before evaluating this proposal, I introduce a second version of the ambiguity thesis.

The second and most recent version of the ambiguity thesis has been defended by Leslie (2015). Her account is unique in that Leslie holds that generics are typically *not* ambiguous. Descriptive generics like (1) and (2), for instance, are said to unambiguously express a single generalization. Still, Leslie agrees that generics with both a descriptive and normative reading – like (5) and (6) – are ambiguous, because, she argues, the kind terms in these sentences are *themselves* ambiguous between a descriptive and normative sense.[7] This second version of the ambiguity view can be represented as follows:

(11) Gen x [Tiger(x)][Striped(x)]

(12) a. Gen x [Philosopher$_1$(x)][Cares-about-the-truth(x)]

b. Gen x [Philosopher$_2$(x)][Cares-about-the-truth(x)]

Leslie argues that words like 'philosopher' and 'man' are ambiguous between a descriptive and a normative sense because the corresponding concepts have a dual character (Knobe et al., 2013; Reuter, 2019). The concepts PHILOSOPHER and MAN are dual character concepts because they encode not only a descriptive dimension consisting of properties typical of kind members, but also a normative dimension consisting of a primary function/social role expected of kind members. Both dimensions are sufficient for categorization. Someone could be considered a philosopher by virtue of her instantiating many of the properties typical of philosophers, but alternatively, also because she performs the primary function/social role of a philosopher. Due to this dual character conceptualization, Leslie argues, words like 'philosopher' and 'man' are ambiguous and can also be

[7] More precisely, Leslie maintains that kind terms like 'philosopher' and 'man' are lexically *polysemous*. Polysemy is a type of lexical ambiguity whereby a single word has multiple senses that are distinct yet related. A typical example of polysemy is the word 'book' which has both a *physical object* sense and an *information* sense (Löhr, 2021). For an overview of current theoretical and experimental issues regarding polysemy, see Falkum & Vicente (2015). Even though Leslie herself uses the term 'polysemy,' I will consistently use the broader term 'ambiguity' to avoid confusion about my objection.

used to denote specifically those people who carry out the primary function/social role associated with the kind. The word 'philosopher,' for example, can be used to denote specifically those people who perform the function/social role conceptually associated with being a philosopher, irrespective of whether they have any of the other properties typical of philosophers.

The distinct functional sense of words like 'philosopher' and 'man' has normative force on Leslie's account, due to what she argues is a *prima facie* social obligation for individuals who satisfy the description of a kind to also carry out the primary function/social role attributed to their kind. Hence the functional-normative sense of 'philosopher' really denotes everyone who exemplifies the *ideal* of a philosopher. The ambiguity of descriptive/normative generics like (5) and (6) can therefore also be represented as follows:

(13) a. Gen x [satisfies-description-of-a-philosopher(x)][cares-about-the-truth(x)]
b. Gen x [exemplifies-the-ideal-of-a-philosopher(x)][cares-about-the-truth(x)]
(14) a. Gen x [satisfies-description-of-a-man(x)][tough(x)]
b. Gen x [exemplifies-the-ideal-of-a-man(x)][tough(x)]

On Leslie's version of the ambiguity thesis, the fact that generics like (5) and (6) have both a descriptive and a normative reading is therefore not really a *generic* phenomenon at all. The two readings that are standardly available for these generic sentences are rather an example of a much broader semantic phenomenon: the ambiguity of kind terms such as 'philosopher' and 'man.' As such, it is unnecessary to say more here about Leslie's otherwise very interesting views on the truth-conditions of *Gen* (see Leslie 2007, 2008). In the next two sections, I argue that neither of the two versions of the ambiguity thesis I have now presented can successfully explain the meaning of descriptive/normative generics.

1.2 The Contradiction Test

One common way to test whether a sentence is ambiguous is to check for the lack of a contradiction in the conjunction of the sentence with its negation (Sennett, 2016; Quine, 1960). In this section, I argue that this so-called 'contradiction test' provides strong evidence against the ambiguity thesis, in particular against the first version of this thesis as defended by Cohen. Since the test is less decisive with regards to Leslie's second version of the ambiguity thesis, further data that weigh against her proposal are provided in the next section.

If a sentence is ambiguous and can be used to express several different propositions, a non-contradictory reading of such a conjunction will be available. For example, since the sentence "The

chicken is ready to eat" is ambiguous, it is possible to read the following conjunction in such a way that it is not contradictory:

(15) The chicken is ready to eat but it's not ready to eat; (we need to cook it first.)[8]

Even though the phrasing is awkward, one can definitely read this sentence in such a way that both conjuncts are true. On the ambiguity account of descriptive/normative generics, these generic sentences are similarly ambiguous between expressing a descriptive generalization and an alternative normative generalization. If that were indeed the case, there should also be a non-contradictory reading available for sentences like (16) and (17) below:

(16) *Philosophers care about the truth but philosophers don't care about the truth.

(17) *Men are tough but men aren't tough.

Against the prediction of the ambiguity thesis, no non-contradictory reading is available for these sentences. There is no reading available according to which one of the conjuncts in these sentences expresses a descriptive generalization whereas the other conjunct states a normative one. The fact that no such non-contradictory reading is available constitutes a decisive problem for the ambiguity thesis, in particular for the first version of this thesis.

If Cohen's view were correct, after all, and generics were syntactically ambiguous between expressing a probabilistic generalization and a rule-based predication, one would expect (16) to have a reading on which it says that *there is a social rule in effect according to which men are tough but any particular man is probably not tough.* This reading on which both conjuncts express an entirely different type of proposition is not available for (16).

Note, furthermore, that on the first version of the ambiguity thesis, *all* bare plural generics are ambiguous, including descriptive ones like (1) and (2). For a descriptive generic like (1), for instance, Cohen's account entails that there should be a statistical reading on which the generic sentence says that any tiger is likely to be striped and an (also descriptive) rule-based reading on which it says that there is a biological rule in effect according to which tigers are striped. However, just like generics with a distinct normative reading, descriptive generics do not pass the contradiction test, as can be seen from the following examples:

(18) *Tigers are striped but tigers are not striped.

(19) *Ravens are black but ravens are not black.

There is no reading of (18) on which the two generics are compatible, nor is there such a reading of (19). Just like generics with both a descriptive and a normative reading, descriptive generics fail

[8] Example adapted from Sennett (2016).

the contradiction test. This is strong evidence against the first version of the ambiguity thesis according to which all generics of the form "*Ks* are *F*" are ambiguous between two different types of generalizations.

The fact that generics about philosophers and men fail the contradiction test, is also a problem for Leslie's ambiguity thesis. On Leslie's version of this thesis, one would expect the conjunctions in (16) and (17) to have a non-contradictory reading because kind terms like 'philosopher' and 'man' can be used to denote two different (not necessarily overlapping) sets of people; one satisfying the descriptive properties of the kind and one exemplifying the ideal of the kind.[9] There should be no contradiction in attributing a property to one of these sets while also denying that it can be attributed to the other set. Yet it appears one cannot just interpret 'philosopher' and 'man' in two distinct ways to resolve the contradiction in (16) or (17), weighing against Leslie's version of the ambiguity thesis.

Still, this contradiction test is less decisive about this version of the ambiguity thesis. After all, despite (16) and (17) being contradictory, there *is* a non-contradictory reading of opposing generics available when special focus is placed on one of the kind terms, like so:

(20) Philosophers don't care about the truth but $[\text{philosophers}]_F$ care about the truth.

(21) Men aren't tough but $[\text{men}]_F$ are tough.

When special focus is placed on one token of the kind term, conjunctions of opposing generics like (20) and (21) are not contradictory. Hence Leslie might well argue that descriptive/normative generics actually pass the contradiction test, supporting her view that these two readings are due to the ambiguity of the kind terms themselves. It is quite plausible, after all, that (20) and (21) are non-contradictory because the focused kind terms denote specifically those people who exemplify the ideals conceptually associated with these kinds.

Yet as I argue in the next section, the fact that kind terms like 'philosopher' and 'man' *can* be used with a functional-normative sense when special focus is placed on them – like in (20) and (21) – does not really show that these terms are ambiguous in the way required to explain the two readings standardly available for descriptive/normative generics.

[9] Not everyone would agree that a non-contradiction test is a suitable way to test for lexical polysemy (Zwicky & Sadock, 1975). However, since Leslie's main argument in favor of her polysemy view is precisely that it is consistent to hold two opposing generics like "boys cry" and "boys don't cry," the non-contradiction test must be a suitable way to test her account even by her own lights.

1.3 Forcing Contexts

The explanandum in this paper is the fact that both a descriptive and a normative reading are readily available for generics such as (5) and (6), even in a default context. There is no further contextual scene-setting required to understand that "Philosophers care about the truth," for instance, can be interpreted in a prescriptive manner as well. If the meaning of the kind terms themselves is to explain this, these kind terms would have to be ambiguous in a very strong way. It would have to be the case that both a descriptive and a functional-normative sense for the kind terms are stored in our lexicon and available to a reader whenever the term is used. Only when both senses of the kind terms are readily available to a reader in this way, could this explain why generics such as (5) and (6) have both a descriptive and normative reading without further contextual cues.

In this section, I argue that kind terms like 'philosopher' and 'man' are *not* ambiguous in this strong way. Leslie provides several examples in support of her view that terms like 'philosopher' and 'man' *can* also have a functional-normative sense, but crucially, these are all examples in which further contextual factors force the reader to look for the non-ordinary use of the kind term intended by the speaker. Once we remove these forcing contextual factors from the examples, the kind terms only have their ordinary sense (which need not be *either* descriptive *or* normative but may well be an underspecified or overspecified combination of both dimensions). As such, these kind terms are not ambiguous in the strong way that could explain why a distinct normative reading is available for some generics even in a default context.

A first type of examples are what Leslie calls 'normatively shifted predications.' These are sentences in which kind terms are used predicatively but appear to predicate that an individual instantiates the idealized social role associated with a kind. Here are two examples:

(22) Hillary Clinton is the only man in the Obama administration.

(23) My sister is always looking for answers to deep questions. Even though she is not a philosopher, if you think about what it really means to be a philosopher, she is a philosopher.

Leslie takes examples like (22) and (23) as evidence for her view that kind terms like 'philosopher' and 'man' can be used to express a distinct social role sense. I will not dispute this claim here. Still, these examples do *not* show that 'philosopher' and 'man' are ambiguous in the strong way required to explain the two readings standardly available for descriptive/normative generics. After all, in both (22) and (23) the context forces the reader to look for the non-ordinary way in which the speaker is using the kind terms. In (22), the speaker flouts the maxim of quality by stating something that is obviously false when the kind term is interpreted in an ordinary manner. In (23), the speaker

instructs the reader to look for 'the real meaning' of the kind term and to thereby avoid the contradiction.

In fact, examples very similar to (22) and (23) show that kind terms like 'philosopher,' 'man,' and 'woman' are not ambiguous in the required strong way. When these kind terms are used predicatively without any further contextual factors forcing an alternative interpretation of them, the sentences in which they appear do not have two separate readings:

(24) Hillary Clinton is the only woman in the Obama administration.

(25) Nietzsche is a philosopher.

Neither (24) nor (25) is ambiguous between two different readings. Both these sentences have only a single reading based on the ordinary meaning of the kind terms. Examples such as these provide evidence against Leslie's claim that social kind terms such as 'woman,' 'man,' and 'philosopher' are ambiguous in a way that could explain the two readings that are by default available for descriptive/normative generics. After all, if the use of these kind terms does not cause sentences like (24) and (25) to have two different readings, why would it cause generics like (5) and (6) to have two different readings? There must be something about the generic formulation as such that is responsible for the two readings of these generic sentences.

For a second type of example, consider sentences in which 'philosophers' and 'men' are used to restrict the scope of a quantifier. Leslie provides several examples in which explicitly quantified generalization featuring these kind terms have a normative reading. Although the following sentences are different from the ones she discusses, they work equally well:

(26) Most [philosophers]$_F$ care about the truth.

(27) All [men]$_F$ are tough.

When special focus is placed on the kind terms, these quantified generalizations have normative force. Again, however, this functional-normative use of the kind term is the result of the context forcing an alternative reading, in this case the interaction of focus with the meaning of the kind term. Without such a focused reading (and without further pragmatic factors), explicitly quantified generalizations only have their ordinary descriptive meaning. This is shown by sentences (3) and (4) mentioned before and repeated here:

(3) Most philosophers care about the truth.

(4) All men are tough.

Neither of these sentences is ambiguous between a descriptive and a normative reading. Yet if Leslie were correct in proposing that kind terms like 'philosopher' and 'man' are ambiguous in a way

that could explain the standard availability of two readings for normative generics, one would expect these quantified generalizations to standardly have two different readings just as well. One would expect (3), for example, to not only have a descriptive reading but also have a reading on which it says that *most people who exemplify the ideal of a philosopher care about the truth*. Similarly, there should be a reading available for (4) according to which it says that *all people who exemplify the ideal of a man are tough*. Yet without a forcing context, these normative readings are not available. The fact that there are no normative readings readily available for these quantified sentences also suggests that it is not the kind terms but the generic formulation of (5) and (6) that is responsible for their two readings.

Taking stock of the paper so far, I have now provided an objection to both versions of the ambiguity thesis. In this section, I have provided an objection against Leslie's view that generics like (5) and (6) have two readings due to the descriptive/normative ambiguity of the kind terms themselves. Even though it may be the case that there is a somewhat conventionalized alternative use of kind terms like 'philosopher' and 'man,' this use of them requires a context that forces a non-ordinary interpretation. Given that no such forcing context is required for generics like (5) and (6) to have both a descriptive and a normative reading, there must be something about the generic formulation itself that is causing these sentences to have two readings. Yet as I have argued in the previous section, it is also not the case that a generic formulation is itself syntactically ambiguous between two different types of generalization. Generics fail the contradiction test (when there is no focused reading of one kind term to force a change of subject.)

Hence an alternative theory of normative generics is required; one according to which their two readings are not a case of syntactical or lexical ambiguity. Such a theory would be successful if it can explain why an additional normative reading is more readily available for some generics than for others and why generics are particularly prone to having this normative reading compared to explicitly quantified generalizations. The contextualist theory defended in the second part of this paper provides these explanations.

3. Descriptive/Normative Generics are Contextually Underspecified

3.1 A Relational Theory of Kindhood

Descriptive/normative generics like (5) and (6) have two different readings in a default context not because they are ambiguous but because they are contextually underspecified. To argue for this position, I will introduce a novel semantic theory for generics, according to which they express context-dependent kindhood generalizations. Other philosophers and linguists have already proposed several interesting semantic theories according to which the content of generics is context-dependent in some way (Greenberg, 2002; Krifka et al., 1995; Nguyen, 2020; Nickel, 2010, 2016;

Sterken, 2015).[10] Discussing these alternative contextualist theories is beyond the scope of this paper. Instead, I will present the essentials of my own theory and argue that it explains the two readings of descriptive/normative generics better than either of the two versions of the ambiguity thesis. I trust that the argument presented here will also be of interest to those who have defended alternative contextualist theories.

Whereas it is quite common to point out that generics characterize the members of a kind, the importance of the notion of *kindhood* for the meaning of generics has not been sufficiently recognized (but see Liebesman & Sterken, 2021). Generic generalizations, I propose, say something about the properties that make a category of individuals into a kind. Speaking somewhat impressionistically first, a generic like "Tigers are striped" says that *being striped* is part of what makes the category *tigers* a kind; it says that this property is co-constitutive of the kindhood of *tigers*. Note already that a property can be part of what constitutes the kindhood of a category without being necessary for membership in the category. *Being striped*, for example, is not necessary for being a member of the category *tigers.* It is, however, a property that individual tigers instantiate in such a way that it is part of what makes the category a kind. I say more about this notion of *kindhood* in the current section, before explaining how it features in the semantics of generic sentences in the next.

Some categories have members that are just a collection of individuals, like say the *ravens in my backyard right now*. Whereas this category can certainly have members, it is not a kind. Other categories, however, have members that are – in some sense to be specified – of the same kind, like the category *ravens* itself. A long and venerable philosophical tradition has aimed to analyze the nature of this difference. What is it that makes some categories of individuals *kinds*, wherein lies their *kindhood*?

One prominent view among philosophers, even if not always phrased in this way, is that the kindhood of a category is grounded in the dependency relation that exists between a set of properties instantiated by members of the category. Some philosophers, for example, have argued that kindhood is grounded in the stable clustering of several properties of category-members, which is one type of dependency relation (Slater, 2015). On this view, the category *philosopher* is a kind, for

[10] The contextualist theory presented here is most obviously inspired by Nickel's account (2010, 2016), according to which *Gen* has the meaning of a universal quantifier restricted in scope to *normal* individuals. This notion of normality, Nickel further argues, is context-dependent and always interpreted in relation to a contextually selected *causal-explanatory strategy*. The theory I present here can be considered an extension of this view; rather than being restricted to normal individuals, generics are restricted based on the notion of kindhood. Kindhood, furthermore, is always interpreted in relation to a *categorizing strategy*, which includes but is not exhausted by explanatory strategies. For more on this, see my (SUPRESSED).

instance, because several properties instantiated by individual philosophers form a stable cluster of generally co-occurring properties. If someone instantiates the properties *writes philosophical texts*, *reads philosophical texts*, and *teaches at a university,* for example, that individual will also tend to instantiate several other properties associated with being a philosopher.

On a second view, the kindhood of a category is rather grounded in the causal relations that exist between several properties of category-members, which is a second type of dependency relation. One version of this causal account proposes that kinds should be thought of as causal networks of properties, where the instantiation of some primary set of causal properties (or a common causal process) is responsible for the instantiation of several secondary ones (Khalidi, 2018). The kindhood of the category *tigers*, for example, could be grounded in the causal dependency relations that exist between several properties of tigers, by virtue of their resulting from the same causal history. The presence of *four legs, tails,* and *stripes* in the population of tigers is ultimately a result of the same causal (selection) process, as is the case for many other properties instantiated by tigers (Godman & Papineau, 2020; Nickel, 2010).

Finally, there are also accounts on which kindhood can be grounded in the functional relations between various properties, another type of dependency (Weiskopf, 2011). In at least some cases, individuals are members of a kind because they share a function that requires them to instantiate several other properties to adequately carry out this function. The category *hearts,* for example, can be considered a kind by virtue of hearts sharing a function that requires the instantiation of several other properties to be carried out.

So, philosophers have proposed at least three different types of dependency relations that can ground the kindhood of a category: statistical cluster dependency, causal dependency, and functional dependency. Often a combination of these three dependency relations is proposed to fully account for the kindhood of a particular kind. In the case of biological species, for example, a causal theory of kindhood is often combined with a statistical cluster view, resulting in a view on which a species' kindhood consists in the reliable clustering of properties due to their sharing a common causal history (Godman & Papineau, 2020; Godman et al., 2020).

Although I cannot fully defend this view here (but see SUPRESSED), one reason why several different analyses of kindhood have been proposed, is that *kindhood* is a relational property of categories. No category is a kind *per se*. Especially in the context of discussing scientific kind categories, several philosophers have now proposed that scientific categories are only kinds *in relation to* the epistemic targets and practices of a particular scientific discipline (Boyd, 2000; Slater, 2013, 2015). These epistemic targets include the set of phenomena one aims to describe, predict, explain, and so forth. The epistemic practices include the set of representational tools used to to describe, the type of explanations that are considered, and so forth. It is in relation to such a set of

epistemic targets and practices that a scientific category can be a kind. A scientific category can be a kind *for* the epistemic targets and practices of particle physics, for example, or *for* those of evolutionary biology.

What then is this relational property of kindhood? A category is a kind for a scientific discipline if the category *accommodates* the epistemic targets and practices distinctive of this discipline, based on the dependency relations that exist between the properties of category-members (Boyd, 2000). The category *proton* is a kind for particle physics, for example, because there is a cluster of properties that always and lawfully co-occurs in protons, due to which this category accommodates the epistemic targets and practices of particle physics. Since epistemic targets and practices vary between scientific disciplines, however, the relativity of kindhood also entails that the type of dependency relations that ground a category's kindhood can vary (Slater, 2013, 2015). Whereas in some disciplines, kindhood requires the exceptionless clustering of several properties (e.g., particle physics), in other disciplines it requires only a looser level of clustering (e.g., sociology), or only that a group of properties can all be explained as the effects of the same causal history (e.g., evolutionary biology).

The relational nature of kindhood generalizes to folk categories as well, resulting in even more ways in which kindhood can be realized. There are many other reasons for categorizing individuals than just the epistemic targets and practices of science, like for example keeping track of the normative expectations that apply to a group of people (Mallon, 2016). More generally then, a category can be a kind in relation to what I call a 'categorizing strategy.' A categorizing strategy is a set of targets and practices with respect to which a category is recognized and used. It is the motivating *rationale* behind a categorization. These categorizing strategies include the epistemic targets and practices distinctive of a scientific discipline, but also the epistemic and more practical concerns that play a role in the recognition and use of other (folk) categories. To give a simple example; one such strategy is to group objects that one can expect to be useful to perform a particular task. A category like *hammers* is a kind in relation to this categorizing strategy; it is a kind of tool. As such, the kindhood of hammers lies in the clustering of a set of properties – like having a handle and a firm head – based upon which hammers fulfil their defining function.

Kindhood can now be defined as follows:

> *Definition Kindhood:* The dependency relation *R* that exists between a set of properties instantiated by members of category *C* is constitutive of the kindhood of *C* relative to categorizing strategy *S*, iff category *C* accommodates strategy *S* by virtue of dependence relation *R*.

In other words, the kindhood of a category relative to a categorizing strategy lies in the dependency relations that exist between several properties of category-members based upon which the category accommodates this categorizing strategy. Each of the properties that stands in this suitable dependency relation, furthermore, is thereby *co-constitutive* of the kindhood of the category. The property *having a handle*, for example, is co-constitutive of the kindhood of *hammers* in relation to the categorizing strategy behind tool categories. This property of *having a handle* stands in the right type of dependency relation to other properties of hammers to accommodate the functional *rationale* behind tool categories.

I have now introduced a relational theory of kindhood (i.e., a category accommodating a categorizing strategy) and a further derivative account of what it is for a property to be co-constitutive of this kindhood (i.e., standing in a suitable dependency relation to other properties of category-members to accommodate a categorizing strategy). In what follows, I support this view on the relational nature of kindhood by arguing that in combination with a kindhood semantics for generics, it explains both the context-sensitivity of generic sentences and the fact that both a descriptive and normative reading is readily available for some of them.

3.2 Contextualist Kindhood Semantics

Based upon the relational account of kindhood proposed in the previous section, we can return to explaining the meaning of generic sentences. Here is a first – still incomplete – account of the truth-condition of generic sentences:

> *Truth-Condition 1:* A generic of the form '*Ks* are *F*' is true iff all *Ks* that – with respect to a determinable of *F* - instantiate a property that is co-constitutive of the kindhood of *Ks*, instantiate *F.*

On this account of the semantics of generic sentences, they express kindhood generalizations. The unpronounced *Gen* operator has the meaning of a universal quantifier but is restricted in scope based on the notion of kindhood. A generic like (1), for instance, says that all tigers that – with respect to their fur pattern – instantiate a property that is co-constitutive of the kindhood of tigers, are striped. In this example, 'fur pattern' is the determinable of *F* because having stripes is one determinate version of a fur pattern.

Since the notion of kindhood only makes sense *in relation to* a categorizing strategy, however, generics are context-dependent. Recall that for a category to be a kind is for this category to accommodate a particular categorizing strategy. Hence for a generic generalization to be truth-evaluable, requires a contextually supplied categorizing strategy. It is only in relation to such a

strategy – that is, in relation to an answer to the question "why categorize this way"? – that it makes sense to consider the properties that are co-constitutive of the kindhood of the category.

The truth-condition of generic sentences can now be spelled out more completely, recognizing their context-sensitivity as follows:

> Truth-condition 2: A generic of the form '*Ks* are *F*' is true in context *C* iff all *Ks* that – with respect to a determinable of *F* – instantiate a property that is co-constitutive of the kindhood of *K* relative to categorizing strategy *S* selected by *C*, instantiate *F.*

On this account of generic meaning then, the generalizing content of a generic sentence depends on a contextually selected categorizing strategy, in relation to which the kindhood of the category is to be understood. How exactly a categorizing strategy is contextually selected is a large question that cannot be fully answered here. Generally, however, it is the intentions of the speaker that select a categorizing strategy from those that are available in the common ground.[11]

In some cases, there is only one possible understanding of a category's kindhood standardly available in the common ground, given how the category and property under discussion are conceptualized. It is common ground, for example, that species categories are kinds in relation to the strategy to group individuals such that their shared history causally explains the reliable clustering of several of their properties. Given this common ground, a speaker who says that "Tigers are striped" will standardly be interpreted as intending to talk about the kindhood of tigers in relation to this categorizing strategy. Hence in a default context, this generic only has a single descriptive reading.

Other generics also have a normative reading, however, because there is also a distinct social role understanding of the category's kindhood available in the common ground, due to the dual character conceptualization of the kind under discussion. Take "Philosophers care about the truth" (i.e., (5)) as a first example. What makes philosophers a kind? That depends. On the one hand, one can conceive of the category *philosophers* as a kind in relation to the categorizing strategy to group people based on their profession, expecting (in the epistemic sense) people with the same profession to share a host of other properties as well. In that case, the kindhood of *philosophers* is constituted by the set of properties that tend to be co-instantiated by philosophers due to these properties being involved in being a professional philosopher. On the other hand, one can also conceive of the category *philosophers* as a kind in relation to the strategy to group people based on their social role. This

[11] More specifically, I believe King's coordination account for demonstratives applies to the context-sensitivity of generics as well. On this account, "the semantic value of a use of a demonstrative d in a context c is that object o that meets the following two conditions: (1) the speaker intends o to be the value of d in c; and (2) a competent, attentive, reasonable hearer who knows the common ground of the conversation at the time of the utterance would know that the speaker intends o to be the value of d in c" (King, 2012, 102). For a similar application of King's coordination account to the context-sensitivity of generics, see Sterken (2015).

results in a grouping of people that must instantiate a particular set of properties if they are to adequately perform their social role. The kindhood of the category *philosophers* in relation to this second categorizing strategy lies in the set of properties upon which carrying out the social role of a philosopher is functionally dependent.

The content of an utterance of (5) depends on the intention of the speaker to either talk about *philosophers* as a professional category or as a social role category. If the first strategy is contextually selected, (5) says that *caring about the truth* is part of the cluster of properties that tend to be co-instantiated due to their involvement in being a professional philosopher. Or more accurately, (5) says that all philosophers that – with respect to their relation to truth – instantiate a property that is co-constitutive of the kindhood of *philosophers* in this way, care about the truth. If, on the other hand, (5) is uttered in relation to the second categorizing strategy, it says: all philosophers that – with respect to their relation to truth – instantiate a property that is part of the set of properties required to adequately carry out the social role of a philosopher, care about the truth.

This second reading has normative force because the notion of a social role is *inherently normative*. As Leslie already said, there is a *prima facie* obligation for members of a kind to carry out the social role attributed to their kind. The normative reading of (5) can therefore also be spelled out as follows: all philosophers that – with respect to their relation to truth – instantiate a property that is part of the set of properties required to adequately carry out the idealized social role of a philosopher, care about the truth.

These same points apply to "Men are tough" (i.e., (6)). Given the dual character conceptualization of the category *men*, this category can be considered a kind in relation to two different categorizing strategies. One strategy is to group people based on an explanatory and predictively fruitful biological property; another strategy is to group people based on their idealized social role. Depending on whether (6) is interpreted in relation to the first or the second strategy, it states a statistical-causal generalization or a normative social role generalization.

In support of this contextualist view on generics, I will present two arguments. Although there are no universally agreed-upon tests for context-sensitivity, the two tests used here set a very high bar (Cappelen & Lepore, 2003). Hence if (descriptive/normative) generics pass these tests, we have good reason to believe their content is context-dependent.

One way to test whether a type of expression is context-dependent, is to check whether it passes the *Inter-Contextual Disquotation Test* (Cappelen & Lepore, 2003). To do so, one checks whether the expression is such that it can be used in a sentence S so that one can truthfully utter an instance of the following schema:

(ICD) Even though it is not the case that S, there can be true utterances of 'S.'

Consider, for example, that an indexical like ‘she’ clearly passes this test. Imagine a context in which one is pointing to a woman who does not speak French. In that context, one can truthfully assert:

(28) Even though she doesn’t speak French, there can be true utterances of ‘She speaks French.’

One can truthfully assert (28) in a context in which one is using ‘she’ to refer to someone who does not speak French, because one readily recognizes that in an alternative context of utterance ‘she’ can refer to a different person that does speak French. If generics are similarly context-dependent, there should also be contexts in which one can truthfully assert an instance of a similar schema. Here is one such case:

(29) Philosophers don’t care about the truth. If you look around, it’s obvious that a vast majority of philosophers don’t actually care about the truth. Of course, one could still say that “Philosophers care about the truth” and say something that is true, but only if one were talking about the properties needed to instantiate the ideal of a philosopher.

In this example, the falsehood of the generic is asserted in relation to a descriptive-statistical categorizing strategy. Yet even within that context, it is possible for the speaker to recognize that an utterance of the generic could also be true in an alternative context, if in that alternative context the operative categorizing strategy is to group people according to their idealized social role. This case shows that it is possible to both assert the falsehood of a generic while at the same time saying that there are true utterances of the same generic sentence in a different context. As such, generics pass the ICD-test.

A second way to check whether a type of expression is context-dependent, is to test whether it blocks so-called *Inter-Contextual Disquotational Indirect Reports* (the IDI-test). Cappelen and Lepore provide the following example to explain this test:

> Take an obviously context sensitive expression, e.g., the first person pronoun "I". Consider an utterance of the sentence "I went to Paris" by Rupert. If Lepore tries to report what Rupert said with "Rupert said that I went to Paris," his report is false because the expression "I" fails to pick out what "I" picked out in the original utterance. The presence of "I" in the disquotational report figures prominently in an explanation of why the report is false. (Cappelen & Lepore, 2003, 23)

Generics similarly block disquotational indirect reports when the context of the original utterance and the context of the report differ with respect to the operative categorizing strategy. Consider, for example, a case in which a conservative father is talking to his children at the dinner table, aiming to

explain to them how men and women are to behave in order to perform their proper role in society. In that context, the father says: “Most men nowadays are weak. This is all wrong. Men are tough.” The next day one of the children is asked by his teacher whether he can list any properties men tend to have. In response to that question, the child reports: “My dad said that men are tough.” This report is false. The original normative ‘social role’ content of the generic utterance does not carry over to the new context, in which the same sentence now expresses a descriptive generalization. As such, generics also pass the IDI-test for context-sensitivity.

Both preceding tests support the view that the content of (descriptive/normative) generics is context-dependent. This context-sensitivity explains why generics are particularly prone to having a normative reading compared to explicitly quantified generalizations; generics *require* a contextually supplied categorizing strategy in relation to which the kindhood of the category is to be understood. If an idealized social role understanding of the category’s kindhood is contextually selected, the generic has normative force. Furthermore, such a normative reading is more readily available for some generics than for others, depending on how the kind is conceptualized. Given that descriptive/normative generics like (5) and (6) concern kinds that are conceptualized in a dual character way, two different categorizing strategies are by default available in the common ground. In this way, the context-sensitivity of generics provides an alternative explanation of how dual character concepts can lead to some generic sentences having both a descriptive and a normative reading even in a default context; these sentences are contextually underspecified without further information about which categorizing strategy the speaker intends to be the operative one.

3.3 Underspecification, not Lexical Ambiguity

In this final section, let me explicitly compare my contextualist view on descriptive/normative generics to Leslie’s ambiguity view. Recall that Leslie proposes that the dual character conceptualization of kinds like *philosophers* and *men* causes the terms ‘philosopher’ and ‘man’ to be ambiguous. Due to ambiguity of these words, generics like (5) and (6) would consequently also be ambiguous between expressing a descriptive and a normative generalization. In the first part of the paper, I objected that kind terms such as ‘philosopher’ and ‘man’ are not ambiguous in the strong way required to explain why both a descriptive and a normative reading are available for (5) and (6) even in a default context, since these terms only have a distinct functional-normative sense in contexts that force a non-ordinary reading.

On a contextualist view, descriptive/normative generics rather have two different readings because kinds like *philosophers* and *men* are conceptualized as having members that share a cluster of properties *and* that are expected to carry out the same idealized social role, requiring them to

instantiate various properties. As such, it is common ground that there are two different ways of conceiving of the kindhood of these categories. Since generics express kindhood generalizations, their content is determined by which of these understandings of kindhood is contextually selected (or more precisely, which kindhood-determining categorizing strategy is selected). In a default context, no further contextual information is available than what is in the common ground, and hence two different readings for the generic sentence are available. That is to say, a reader of (5) or (6) recognizes the two possible ways of understanding these sentences not because they recognize that kind terms like 'philosopher' and 'man' are ambiguous, but because they recognize that there are two different ways of understanding what it is that makes philosophers and men a *kind*; a shared cluster of properties or a shared idealized social role.

This underspecification account of descriptive/normative generics also explains the two phenomena that prompted Leslie's ambiguity view, without facing the objections levelled against this view in the first part of the paper. One phenomenon Leslie refers to in support of her view is that "there is a sense in which one can coherently hold both that *boys don't cry* and that *boys do cry*" (Leslie, 2015, 113). This is true but it is not a result of the ambiguity of 'boys.' Rather, it is a result of the context-sensitivity of *Gen.* One can consistently hold two opposing generics *provided they are held in relation to two different categorizing strategies*. One can consistently hold that "boys cry" and that "boys don't cry" since both sentences can be interpreted in relation to a different categorizing strategy. Yet it is contradictory to hold that "Boys don't cry but boys cry" since this sentence is interpreted in relation to a single categorizing strategy. Hence why descriptive/normative generics fail the contradiction test (without a focused reading of one kind term changing the subject).

A second phenomenon referred to by Leslie (2015) in support of her ambiguity view, is that normatively shifted predications are often argumentatively supported by a generic. For example, "Hillary Clinton is the only man in the Obama administration" might be supported by the speaker adding that "Men are tough." Leslie explains this by arguing that there is a semantic connection between the generic sentence and the normative predication; in both cases the kind term is used to denote specifically those people who instantiate the social role of the kind. Yet the contextualist view I have defended explains this phenomenon equally well. There is indeed a semantic connection between the premise (i.e., generic) and the conclusion (i.e., predication). Generics are restricted in scope based on a contextually selected understanding of kindhood. If in the context, the kindhood of *men* is understood as being grounded in their idealized social role requiring the instantiation of many other properties, then the generic (roughly) says that all men who instantiate a property that is co-constitutive of their kindhood in this way, are tough. As such, the *scope* of a generic generalization like "Men are tough" that is uttered in the context of a discussion of men's social role, is identical to the denotation of the normatively shifted understanding of 'man.' This semantic connection can

explain why generics are often provided in argumentative support of normatively shifted predications.

Conclusion

Some generics, like "Philosophers care about the truth" and "Men are tough," have both a descriptive and a normative reading in a default context. I have argued that this is not due to any syntactical or lexical ambiguity. Generics fail the contradiction test. Furthermore, even though kind terms like 'philosopher' and 'men' *can* be used to convey a functional-normative sense, this is only in contexts that force a non-ordinary reading. Instead, generics are prone to having different readings because they express kindhood generalizations and because *kindhood* is a relational property. The truth-condition of a generic sentence depends on a contextually selected categorizing strategy, determining which dependency relations ground a category's kindhood. When a generic sentence is uttered in relation to the categorizing strategy to group people based on their idealized social role, it has normative meaning. When a generic has both a descriptive and a normative reading, it is because two different categorizing strategies are contextually available.

References

Boyd, R. (2000). Kinds as the "Workmanship of Men": Realism, Constructivism, and Natural Kinds. In Nida-Rümelin, J. (Ed.), *Rationalität, Realismus, Revision: Vorträge des 3. Internationalen Kongresses der Gesellschaft für Analytische Philosophie*, (pp. 52-89). Berlin: De Gruyter.

Cappelen, H., & Lepore, E. (2003). Context shifting arguments. *Philosophical perspectives*, *17*, 25-50.

Cohen, A. (2001). On the generic use of indefinite singulars. *Journal of semantics*, *18*(3), 183-209.

Del Pinal, G., & Reuter, K. (2017). Dual character concepts in social cognition: Commitments and the normative dimension of conceptual representation. *Cognitive science*, *41*, 477-501.

Falkum, I. L. & Vicente, A. (2015). Polysemy: Current perspectives and approaches (editorial), *Lingua*, 157, 1-16.

Godman, M., & Papineau, D. (2020). Species have historical not intrinsic essences. In A. Bianchi (Ed.), *Language and Reality from a Naturalistic Perspective.* Philosophical Studies 143, (pp. 355-367). Springer Press.

Godman, M., Mallozzi, A., & Papineau, D. (2020). Essential Properties are Super-Explanatory: Taming Metaphysical Modality. *Journal of the American Philosophical Association*, 6(3), 316-334.

Greenberg, Y. (2002). Two kinds of quantificational modalized genericity, and the interpretation of bare plural and indefinite singular NPs. In B. Jackson (Ed.) *Proceedings of SALT. Cornell University: Cornell Linguistics Circle*.

Hesni, S. (2021). Normative generics: Against semantic polysemy. *Thought: A Journal of Philosophy*, *10*(3), 218-225.

Khalidi, M. A. (2018). Natural kinds as nodes in causal networks. *Synthese*, *195*(4), 1379-1396.

Knobe, J., Prasada, S., & Newman, G. E. (2013). Dual character concepts and the normative dimension of conceptual representation. *Cognition*, *127*(2), 242-257.

King, J. (2012). The Metasemantics of Contextual Sensitivity. In A. Burgess & B. Sherman (Eds.), *Metasemantics: New Essays on the Foundations of Meaning*. Oxford: Oxford University Press.

Krifka, M., Pelletier, F. J., Carlson, G. N., ter Meulen, A., Chierchia, G., and Link, G. (1995). "Genericity: An Introduction," in G. N. Carlson and F. J. Pelletier (eds.), *The Generic Book*, Chicago: University of Chicago Press.

Laskowski, N. G. (2020). Moral Constraints on Gender Concepts. *Ethical Theory and Moral Practice*, 1-13.

Leslie, S.J. (2007). Generics and the structure of the mind. *Philosophical perspectives*, *21*, 375-403.

Leslie, S.J. (2008). Generics: Cognition and acquisition. *Philosophical Review*, *117*(1), 1-47.

Leslie, S.J. (2015). "Hillary Clinton is the only man in the Obama administration": Dual Character Concepts, Generics, and Gender. *Analytic Philosophy*, *56*(2), 111-141.

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

Löhr, G. (2021) Does polysemy support radical contextualism? On the relation between minimalism, contextualism and polysemy, *Inquiry*, online first.

Mallon, R. (2016). *The construction of human kinds*. Oxford University Press.

Nguyen, A. (2020). The radical account of bare plural generics. *Philosophical Studies*, *177*(5), 1303-1331.

Nickel, Bernhard. 2010. Ceteris Paribus laws: generics & natural kinds. Philosophers' Imprint 10(6): 1-25.

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

Reuter, K. (2019). Dual character concepts. *Philosophy Compass*, *14*(1), e12557.

Sennet, A. (2016*).* Polysemy*. Oxford handbooks online.*

Slater, M. H. (2013). Are Species Real? An essay on the metaphysics of species. London: Palgrave-Macmillan.

Slater, M. H. (2015). Natural kindness. *The British Journal for the Philosophy of Science*, *66*(2), 375-411.

Sterken, R. (2015). Generics in context. *Philosopher's Imprint, 15*(21), 1–30.

Sterken, R. K. (2017). The meaning of generics. *Philosophy Compass*, *12*(8), e12431.

Weiskopf, D. A. (2011). The functional unity of special science kinds. *The British Journal for the Philosophy of Science*, *62*(2), 233-258.

Quine, W. V., 1960. *Word and Object*, Cambridge, MA: MIT Press.