UNIVERSITA' VITA-SALUTE SAN RAFFAELE

CORSO DI DOTTORATO DI RICERCA IN FILOSOFIA

Curriculum in Philosophical Issues in an International Perspective

SOCIAL GENERICS AND INFERENTIAL ASYMMETRY

Tutore: Chiar.ma Prof.ssa Claudia Bianchi

Co-Tutore: Chiar.ma Prof.ssa Susan A. Gelman

Tesi di DOTTORATO di RICERCA di Federico Cella

matr. 012229

Ciclo di dottorato XXXIII

SSD M-FIL/05

Anno Accademico 2019/2020

Claudia Bianchi

ABSTRACT

Generics (e.g., "Ravens are black") express generalizations over categories or their members. Generic language is the subject of numerous works in linguistics, philosophy of language, and psychology. Previous research (Cimpian et al., 2010; Brandone et al., 2015) showed that generics about animals are interpreted as broadly true of members of a kind, yet also accepted based on minimal evidence. This inferential asymmetry is important for suggesting a mechanism by which unfounded generalizations may flourish; nonetheless, prior research has not examined whether this finding extends to the interpretation of *social generics* (i.e., generics about groups of people). Accordingly, in collaboration with Marchak, Bianchi, and Gelman, I investigated in four studies (n = 665) whether the inferential asymmetry elicited by generics differs as a function of domain (animals vs. people). Before presenting our studies, I consider some of the main theoretical and experimental works carried out on generics.

In Chapter 1, I consider the linguistic issues raised by generics. Genericity comprehends a great variety of linguistic phenomena; furthermore, the semantics of these sentences is not straightforward and resist formal analysis (e.g., Krifka et al., 1995).

In Chapter 2, I discuss some of the main cognitive implications of generics. I consider the hypothesis that generics are default generalizations (e.g., Leslie & Gelman, 2012) and the influence of generic language on social cognition (e.g., Rhodes et al., 2012).

In Chapter 3, I mainly discuss my research. In our studies, participants were randomly assigned to either an *Implied Prevalence* task (given a generic, asked to estimate the prevalence of a property) or a *Truth-Conditions* task (given prevalence information, asked whether a generic was true or false). We found that the inferential asymmetry did not differ across domains and also extended to properties varying in type (physical vs. non-physical) and valence (dangerous vs. neutral). Moreover, there were differences as a function of property valence in the Implied Prevalence task and a small but consistent interaction between domain and prevalence in the Truth-Conditions task. I discuss the implications of these results for the semantics of generics and the relation between generics and stereotyping.

Table of contents

INTRODUCTION 4

1. WHAT ARE GENERICS? 8

1.1 Varieties of generics 8

1.1.1 Preliminary syntactic distinctions 10

1.1.2 Generics with nominal arguments 12

1.1.3 Kind-referring and object-referring NPs 13

1.1.4 Characterizing and particular sentences 20

1.1.5 The focus of my research 24

1.2 Generics versus non-generics 29

1.2.1 Generics versus quantified sentences 29

1.2.2 Generics versus existentials 36

1.3 What do generics mean? 40

1.3.1 Semantic analyses of generics: an overview 41

1.3.2 A probability-based approach 44

1.3.3 A content-based approach 50

2. GENERICS, COGNITION, AND STEREOTYPING 61

2.1 Generics as default generalizations 61

2.1.1 Motivations 62

2.1.2 The hypothesis 68

2.1.3 Confirmations 71

2.2 Generics and social cognition 81

2.2.1 The harmful consequences of social essentialism 82

2.2.2 Generics and social essentialism: a "self-reinforcing cycle" 85

2.2.3 The content-based approach and the pitfalls of social generics 92

2.2.4 Controversies 100

3. SOCIAL GENERICS AND INFERENTIAL ASYMMETRY.....109

3.1 Background issues 109

3.1.1 The inferential asymmetry 110

3.1.2 The role of domain in the interpretation of generics 115

3.1.3 The role of property type and valence in the interpretation of generics 117

3.1.4 The present studies.....119

3.2 Study 1 120

3.2.1 Method 120

3.2.2 Results and discussion 122

3.2.3 Conclusions.....123

3.3 Study 2 124

3.3.1 Method 124

3.3.2 Results and discussion 126

3.3.3 Conclusions.....129

3.4 Study 3 131

3.4.1 Method 132

3.4.2 Results and discussion 134

3.4.3 Conclusions.....137

3.5 Study 3a 140

3.5.1 Method 140

3.5.2 Results and discussion 142

3.6 General discussion 143

3.6.1 The robustness and generalizability of the generic asymmetry.....144

3.6.2 The effect of property type and valence in the interpretation of generics 145

3.6.3 The effect of domain and property valence on generics' acceptability.....147

3.6.4 Implications 149

3.6.5 Conclusions 151

CONCLUDING REMARKS 153

References 156

Webography 173

Appendix 175

Appendix S1 192

Appendix S2 193

Appendix S3 194

Appendix S4 195

INTRODUCTION

Generics are sentences such as "Ducks lay eggs", that express general claims regarding categories or their members and that lack explicit quantifiers. They contrast with non-generic sentences such as "'Most ducks lay eggs" (which is quantified) and "Ducks are laying eggs"/ "This duck lays eggs" (which refer to specific individuals). Generic noun phrases can have different syntactic forms in English, including bare plural (as in "Leopards have spots"), definite singular (as in "The raven is black"), and indefinite singular (as in "A tiger is striped"). Generics have both puzzling semantic features and distinctive cognitive implications; for these reasons, they have attracted the interest of linguists (e.g., Carlson, 1977; Cohen, 1996; Krifka, Pelletier, Carlson, ter Meulen, Chierchia, & Link, 1995), philosophers (e.g., Langton, Haslanger, & Anderson, 2012; Leslie, 2008; Saul, 2017), and psychologists (e.g., Gelman, 2003; Cimpian, Brandone, & Gelman, 2010; Rhodes, Leslie, & Tworek, 2012).

One intriguing finding is that generics about novel animal categories, such as "Morseths have silver fur," are characterized by an inferential asymmetry: they are interpreted as referring to nearly all members of the kind, despite being accepted even if the ascribed property is present in relatively few members (e.g., 10%) (Cimpian et al., 2010; Brandone, Gelman, & Hedglen, 2015). This finding has important implications for several real-world phenomena, which range from how generic language affects learning processes to propaganda in political discourse. The asymmetry is also relevant to stereotyping. Indeed, *social generics* – i.e., generics about groups of people – are a common means to express stereotypes (Gelman, Taylor, & Nguyen 2004), and could be particularly misleading and pernicious in communication if they displayed both strong prevalence implications and loose acceptance conditions. Stereotypes of the generic form "Girls are bad at math" or "Italians are mobsters" could be easily generated on the basis of few instances, yet leading others to believe that almost every member of the relevant category presents the ascribed properties (e.g., Bian & Cimpian, 2017; Brandone et al. 2015; Cimpian et al., 2010; Hammond &

Cimpian, 2017; Langton, Haslanger, & Anderson, 2012; Rosola & Cella, 2020). Yet, to date, it had not been investigated whether the asymmetry would be found in social categories. Accordingly, this dissertation investigates the robustness and generalizability of the inferential asymmetry for social generics.

This research was developed at the Conceptual Development Lab at the University of Michigan through discussion and collaboration with Kristan Marchak (University of Alberta), Claudia Bianchi (Vita-Salute San Raffaele University), and Susan Gelman (University of Michigan). To examine whether the inferential asymmetry found by Cimpian et al. (2010) and Brandone et al. (2015) extends to the interpretation of social generics, we conducted four studies (n = 665). These studies also examined whether the inferential asymmetry varies as a function of property type (physical and non-physical) and property valence (dangerous and neutral) in two domains (animals and social categories). Before focusing on our studies, I discuss some of the main theoretical and empirical works previously carried out on generics. In particular, my dissertation is structured as follows.

In Chapter 1, I consider some issues concerning the linguistic features of generics. From a theoretical point of view, genericity is a rather broad and complicated phenomenon. In the first part, I discuss the varieties of generics and clarify what type of generic sentences are the focus of the present research. The purpose of this technical overview is to circumscribe the manifestations of generic language relevant to my research. In the second part, I compare the linguistic features of generics with the linguistic features of some non-generic constructions that resemble generics, but only in some ways. Finally, in the third part I discuss two semantic theories that have been put forward to account for the meaning of generics.

In Chapter 2, I review the cognitive implications of generics. Specifically, in the first part I consider the hypothesis that generics are default generalizations, together with the empirical data that motive such hypothesis and the experiments that confirm it. In the second part, instead, I consider the relation between generics and social cognition. I specifically focus on the use and interpretation of social generics as stereotypes together with their ability to promote essentialist beliefs (i.e., the belief that the members of a category share an inherent

and non-obvious nature; see Gelman, 2003; Medin & Ortony, 1989; Rhodes et al., 2012). I also consider some controversies that arose within the debate.

Finally, in Chapter 3, I mainly focus on my research. In the first part, I review several key background issues, including the generic inferential asymmetry and potential variation of the asymmetry as a function of domain (animal vs. people), property type (physical vs. non-physical), and property valence (neutral vs. dangerous). In the following parts, instead, I present our studies and discuss our findings, together with their implications for the semantics of generics and for stereotyping.

In our studies, participants were randomly assigned to either the *Implied Prevalence* task or the *Truth-Conditions* task and provided with generics statements about novel animals or people. In the Implied Prevalence task, we measured the average prevalence level implied by generics. In the Truth-Conditions task, instead, we measured the average prevalence levels that participants needed to accept generics. Subsequently, we examined whether the average prevalence level implied by generics is significantly higher than the average prevalence level at which these statements are accepted. In Study 1, we examined whether the generic asymmetry found in animal categories replicates using an improved methodology and a larger online sample. In Study 2, we investigated whether the generic asymmetry varies as a function of domain (animals vs. people) and property type (physical vs. non-physical). In Studies 3 and 3a, we tested whether the generic asymmetry varies as a function of domain (animals vs. people) and property valence (neutral vs. dangerous).

Our studies show that the inferential asymmetry is not significantly different across domains and hold for properties varying in type and valence. However, our findings suggest that the average prevalence level implied by generics is affected by property valence, leading to higher estimates for generics expressing neutral properties (e.g., “Frams are short”) than generics expressing dangerous properties (e.g., “Morseths are violent”). Finally, in the Truth-Conditions task of both Study 2 and Study 3 we observed a small but consistent interaction between domain and prevalence: social generics were more likely than generics about animals to be accepted at the lowest prevalence levels, whereas generics about animals were more likely than social generics to be accepted at the highest prevalence levels.

Overall, these findings deepen our understanding of the semantics of generics of both animal kinds and social groups, together with the relation between generics and stereotyping. Furthermore, they lay the groundwork for future research. Promising directions include investigating whether the generic asymmetry arises in artifact categories and in multiple languages; examining whether generics having different verbs elicit different asymmetries; and further exploring the role of domain and property valence in the interpretation of generics.

1. WHAT ARE GENERICS?

Generics are ubiquitous. We produce and hear instances of them over and over again every day, yet these sentences have been very challenging to analyze. In this chapter, I outline the reasons why generics pose a problem much debated in linguistics and philosophy of language. Since a great variety of phenomena fall under the general label of "generics", I devote the first part to clarifying the kind of generic sentences I focus on in my research. In the second part, I expand on the distinction between this type of generics and other non-generic constructions - that is, *quantified* and *existential* sentences. Finally, in the third part, I discuss two accounts of the generics covered by my research.

1.1 Varieties of generics

The linguistic category of *generics* (from Latin *genus* – i.e., "kind") includes a wide range of different phenomena. Such a variety might generate ambiguities about whether or not certain sentences fall in the present analysis. To avoid this possibility, I start by clarifying what type of generic sentences are the object of my research. My research focuses on people's interpretation of a type of generics known in the literature as *bare plural characterizing sentences*, like "Tigers are striped". Specifically, I am interested in testing whether the inferential asymmetry found by Cimpian, Brandone, and Gelman (2010) and Brandone, Gelman, & Hedglen (2015) with bare plural characterizing sentences about animals extends to bare plural characterizing sentences about social categories. In the next subsections, I explain which sentences fall under this category and which do not. To this end, I mainly rely on the classification proposed by Krifka, Pelletier, Carlson, ter Meulen, Chierchia, and Link in the introduction of *The Generic Book* (1995), one of the most important collective texts in the linguistic literature on generics. I also integrate this traditional classification with the previous and following literature.

Before proceeding, some preliminary clarifications are required. First, in the linguistic literature on generics, it is common to use sentences about animals as examples of generics. In this section, my purpose is classificatory only; thus, in line with the relevant literature, I mainly use sentences about animals. However, it is important to bear in mind that generics can also be used to describe any other member of the biological domain (e.g., plants and bacteria), but also natural objects (e.g., minerals), artifacts (e.g., tools), and – importantly – members of social categories (e.g., women and men).

Second, in the examples provided, I use the diacritic "?" to indicate oddness and "#" to indicate whether a sentence is ill-formed. In the literature on generics, there is a tendency to use the "felicitous/infelicitous" distinction to indicate whether a sentence is constructed according to the syntactic and semantic rules of a language or not (e.g., Carlson 1977; Greenberg 2003; Krifka 1987; Krifka et al. 1995; Leslie 2012a). Notice, though, that this terminology derives from the literature on speech acts. In particular, Austin (1962) talks about "infelicity" in reference to illocutionary acts – not in reference to the locutionary (or perlocutionary) ones. The notion of "infelicity" could be *extended* in order to take into account also locutionary flaws (as a semantic ill-formation of a sentence, but also as a Freudian slip, a pronunciation error, and so on), but – in order to avoid misunderstandings – I consider it more appropriate to describe the kind of failure at stake as "ill-formedness".

Finally, in describing the syntactic variations of generics, I always consider English generics. Indeed, other languages present different syntactic variations, which are not the object of the present research. For example, consider the bare plural characterizing sentence "Tigers have stripes" again. The syntactic form of this sentence is not possible in some other languages; for example, the corresponding bare plural sentence in Italian – i.e., "#Tigri hanno le strisce" is ill-formed. The equivalent plural characterizing sentences in Italian can only have a definite plural syntactic form, as in "Le tigri hanno le strisce" – which in English would be "The tigers are striped". This sentence is well-formed in English as well. However, definite plural sentences in English are rarely interpreted as generics (e.g., Gelman, 2003); rather, they are usually interpreted as sentences concerning specific individuals.

1.1.1 Preliminary syntactic distinctions

With the general label "generics" it is possible to denote two different broad phenomena: a) sentences with *nominal arguments* (e.g., Carlson, 1977; Krifka et al. 1995) and b) sentences with *verbal arguments* (e.g., Carlson, 1977; Conrad, 1982). I first consider generics with nominal arguments, which have been the main focus in the literature on generics.

Generics with nominal arguments concern kinds, individuals, or events. They sometimes express generalizations, while sometimes they concern particular facts about whole kinds. These generics present important syntactic variations, as their subject noun phrases, or NPs, can be either count nouns, bare mass nouns, or proper nouns.[1] (1) offers examples of generics having a count noun NP, which have been the most studied:

(1)

a. Humans are social animals.

b. The parrot nests in tree hollows.

c. An acoustic guitar has a resonating chamber.

Generics containing count noun NPs can have different syntactic forms, including *bare plural* (or BP), *definite singular* (or DS), and *indefinite singular* (or IS). (1a) is an example of a BP generic – that is, a generic having a plural NP and lacking a determiner. (1b) is an example of a DS generic: its subject is singular, and it is preceded by the definite article "the". Finally, (1c) is an example of an IS generic, which contains a singular NP preceded by the indefinite singular article "a". The syntactic forms of generics in (1) can be generalized as in (2), where *K*s stands for a plural subject, *K* a singular subject, and *F* a property:

(2)

a. Ks are/have F.

b. The K is/has F.

[1] It should be noted that some generic sentences can also be sentences whose direct object is a generically interpreted NP, as in "Fleming discovered penicillin". However, these are marginal cases that I do not consider in this work.

c. A K is/has F.

Conversely, generics having bare mass noun NPs, like (3a), and generics having a proper noun NP, like (3b), received less attention:

(3)

a. Water is a polar inorganic compound.

b. Jane goes for a run before breakfast.

The second category of generics includes sentences with verbal arguments, which express generalizations over actions or processes. Such sentences can have different types of verb-based arguments, including *gerundives*, as in (4a), *infinitives*, as in (4b), *nominalizations*, as in (4c), and *derived nominals*, as in (4d):

(4)

a. Swimming is healthy.

b. To swim is to keep fit.

c. Discussions about the environment are good.

d. The resistance is overcoming the crisis.

At this point, I can make a first broad selection. The focus of this research is on the interpretation of generics about social categories and animal kinds. Thus, I can leave the category of generics with verbal arguments behind. From now on, I will use the term "generics" to refer to generics with nominal arguments exclusively. At the end of the present section, the set of sentences indicated by the term "generics" will be further restricted and simplified.

1.1.2 Generics with nominal arguments

In addition to a broad syntactic variation, generics present important semantic differences. Based on the distinction proposed by Krifka (1987; partly in collaboration with Gerstner) and subsequently reworked by Krifka et al. (1995), two distinct phenomena have been categorized as generics in the linguistic literature: a) *sentences having a kind-referring NP* and b) *characterizing sentences*.

Sentences having a kind-referring NP, originally called *Definite generics* by Krifka (1987), concern kinds, like (5):

(5) The dog is a sub-species of the wolf.

Conversely, characterizing sentences, originally called *Indefinite generics* by Krifka (1987), express generalizations like (6a) and (6b):

(6)

a. Dogs bark.

b. A dog barks.

Krifka (1987) labels sentences having a kind-referring NP "Definite generics" because their NPs are often DSs, whereas the NPs of characterizing sentences are often ISs or BPs, hence the label "Indefinite generics". Krifka, et al. (1995) question the use of this terminology because, as they put it, «not only are these not universal properties of the two types; the terms also incorrectly suggest that both types are somehow on a par in that they are manifested by [NPs]» (Krifka et al., 1995: 4, footnote 3). In their view, sentences having a kind-referring NP are generics in virtue of their NP, whereas genericity in characterizing sentences is a feature of the whole sentence.[2] Since the two phenomena are manifested at different levels, they are not in contrast; rather, they sometimes overlap. This is the case of generics

[2] Contrary to this traditional view, Mari, Beyssade, and Del Prete (2013) argue that some characterizing sentences, as for instance "John smokes a cigar after dinner", are generics in virtue of their verbal phrase.

expressing generalizations over kinds. For example, (7) expresses a generalization that holds for the kind *Canis familiaris*:

(7) The dog has four legs.

Although the two phenomena can co-occur, Krifka et al. (1995) argue that sentences having a kind-referring NP and characterizing sentences should be analyzed separately, as they present important linguistic differences and some sentences display one type of genericity but not the other. For example, kind-referring NPs can also occur in *particular* sentences, which concern specific properties, events, or episodes, as in (8):

(8) The dog was domesticated in Eurasia.

(8) does not express a generalization; indeed, it reports a specific episode concerning the kind *Canis familiaris*. Characterizing sentences, instead, can also have *object-referring* NPs, which concern individuals instead of kinds, as in (9):

(9) A dog has a great sense of smell.

For reasons that I discuss in the next subsection, sentences like (9) are often described as expressing generalizations based on properties of individuals instead of properties of kinds (e.g., Krifka et al. 1995). In the next subsection, I consider the functioning of sentences having a kind-referring NP by comparing it with the functioning of sentences having an object-referring NP.

1.1.3 Kind-referring and object-referring NPs

Sentences having a kind-referring NP involve reference to kinds. They are opposed to sentences with an object-referring NP, which involve reference to individuals instead. Examples of sentences having a kind-referring NP are the following:

(10)

a. T-rex lived about 66-68 million years ago.

b. The T-rex was first discovered in Wyoming.

c. T-rex are large.

d. The T-rex is extinct.

e. Amber preserves trapped animals.

Conversely, the examples in (11) are object-referring sentences:

(11)

a. The T-rex was crossing the river.

b. A T-rex has a heavy skull.

c. FMNH PR 2081 is a well-preserved T-rex.[3]

d. This amber is roughly 130 million years old.

Sentences (10a)-(10d) concern the kind *Tyrannosaurus rex*, while (10e) concerns a type of resin. The sentences in (11), instead, involve reference to individual members of the same kinds. Both kind-referring and object referring NPs can be paired with *dynamic* predicates – that is, predicates concerning particular facts or events in which the subject referent is involved. (10a) and (10b) are examples of kind-referring NPs paired with dynamic predicates; (11a) is an example of an object-referring NP paired with the same type of predicate. Both kind-referring and object referring NPs can also be paired with *stative* predicates, which concern stable properties of the subject referent. (10c), (10d), and (10e) are examples of kind-referring NPs paired with stative predicates; (11b), (11c), and (11d), instead, are sentences constituted of an object-referring NP paired with other stative predicates.

[3] FMNH PR 2081, also called "Sue" after its discoverer Sue Hendrickson, is one of the most complete fossilized skeletons of a T-rex.

The kind-referring NPs in (10) have three syntactic forms. The NPs in (10a) and (10c) are BPs. The NPs in (10b) and (10d) are DSs. Finally, the NP in (10e) is a bare mass noun. I did not include a sentence with an IS NP amongst sentences with kind-referring NPs because ISs, with one exception I discuss below, cannot refer to kinds. In the literature on generics, it is widely assumed that IS NPs can only be object-referring (e.g., Carlson, 1977; Jespersen, 1927). This assumption is supported by the observation that IS NPs cannot be combined with a type of stative predicates that Carlson (1977) calls *kind-level predicates*. Kind-level predicates like "extinct", "rare", and "widespread" require a kind as an argument: an individual T-rex can be large or well-preserved, but it cannot be extinct or widespread.[4] As a consequence, an object-referring NP cannot be paired with a kind-level predicate. For example, contrast sentences (12a) and (12b) with sentences (12c) and (12d):

(12)

a. T-rex are extinct.

b. The tiger is widespread.

c. # A T-rex is extinct.

d. # A tiger is widespread.

(12a) and (12b) are perfectly fine, whereas (12c) and (12d) are ill-formed. On the contrary, IS NPs can be paired with other stative predicates, as in (13a) and (13b):

(13)

a. A T-rex is large.

b. A tiger is striped.

The sentences in (13) are both well-formed; however, their NPs can only be interpreted as object-referring.

[4] Notice, though, that Krifka et al. (1995) deviate from the literature on generics in considering "widespread" and "rare" as kind-level predicates on a par with "extinct". In their view, the former can sometimes be combined with non-specific object-referring NPs (for more discussion, see Krifka et al., 1995: 95 – 98).

Is there a way to disambiguate between kind-level and other predicates without combining them with an IS NP? Cohen (2001) proposes that verifying whether the modification of the sentence with an overt adverb of quantification (e.g., *usually, always, sometimes, often*, *never*) is impossible constitute a test for cases of direct kind predications. For example, (10c) does not contain a kind-level predicate and can be reformulated as follow:

(14) T-rex are usually large.

(10d), instead, contains a kind-level predicate; the combination of (10d) with the same adverb of quantification results in (15), which is ill-formed:

(15) # The T-rex is usually extinct.[5]

It is important to note, however, that IS NPs can be kind-referring when they receive a *taxonomic* reading. Taxonomic NPs are kind-referring NPs ranging over sub-kinds (e.g., Galmiche 1985; Declerck, 1987, 1991; Krifka et al. 1995). For example, the family *Felidae* includes, among others, lions, leopards, and tigers; the clade *Dinosauria*, instead, include triceratops, brachiosaurus, and T-rex. The following sentences show that taxonomic IS NPs can be combined with kind-level predicates:

(16)

a. A feline, namely the tiger, is widespread.

b. A dinosaur (the T-rex) is extinct.

(16a) and (16b) are both well-formed. Taxonomic IS NPs can also be combined with dynamic and stative predicates and be interpreted as kind-referring, as shown by the examples in (17):

[5] It should be noted that this test works only for certain kind-level predicates. For example, "Kangaroos are usually rare, but not in Australia" is well-formed. This observation is compatible with Krifka et al.'s (1995) view, according to which "widespread" and "rare" are not kind-level predicates on a par with "extinct" (see the previous footnote).

(17)

a. A dinosaur, namely the T-rex, lived about 66-68 million years ago.

b. A dinosaur (the T-rex) is large.

When sentences like the ones in (16) and (17) are interpreted in a non-taxonomic reading, they are either ill-formed, as (12c) and (12d), or cannot be interpreted as kind-referring, as (13a) and (13b).

A further observation is that, as noted by Krifka et al. (1995), both kind-referring and object-referring NPs can be either *specific* or *non-specific*. Specific NPs, contrary to non-specific NPs, refer to particular entities. For example, consider the following sentences:

(18)

a. The T-rex / A dinosaur, namely the T-rex, has short arms.

b. A dinosaur species can be classified based on several criteria.

c. FMNH PR 2081 / A T-rex, namely FMNH PR 2081, has short arms.

d. A T-rex has short arms.

(18a) has a specific kind-referring NP, whereas (18b) has a non-specific kind-referring NP. Conversely, (18c) has a specific object-referring NP, whereas (18d) has a non-specific object-referring NP.

A few NPs are unambiguously kind-referring. This is the case of bare mass nouns, as in (10e), or NPs with special lexical content, as in (18b) or in the following sentences:

(19)

a. This kind of dinosaur has short arms.

b. This kind of amber is called “blue amber”.

Except for these particular cases, BP and DS NPs can have both a kind-referring and an object-referring reading (e.g., Krifka et al., 1995; Chierchia, 1998).[6] Since they can be interpreted as taxonomic NPs, IS NPs have an ambiguous reading as well. Some linguistic tests can help to disambiguate between kind-referring and object-referring NPs. In what follows, I briefly report two tests that are well-known in the literature.

As already discussed, a first simple way to determine whether a certain type of NP can be used as kind-referring is to pair it with a kind-level predicate. An object-referring NP cannot be paired with a kind-level predicate; thus, if the resulting sentence is ill-formed, that type of NP cannot be used as kind-referring (Krifka et al., 1995). A second way consists of comparing a certain NP's behavior when it occurs in particular sentences that are both true and non-negated. Specifically, some scholars noted that when this type of sentences have an object-referring NP they are *upward-entailing*: they remain true even if the subject-term is replaced by a more inclusive one (e.g., Krifka et al. 1995; Laca, 1990; Lawler, 1973). For instance, assuming the truth of (20a), it is possible to replace the term "tiger" in (20a) with the term "animal" in (20b) salva veritate:

(20)

a. Tigers were in my backyard.

b. Animals were in my backyard.

This does not apply to sentences having a kind-referring NP. While (21a) is true, (21b) is false:

(21)

a. Tigers disappeared from Bali during the last century.

b. Animals disappeared from Bali during the last century.

[6] Other scholars have sometimes considered BP NPs as the plural counterpart of IS NPs (see Burton-Roberts, 1976 and Kratzer, 1995 for discussion); yet others have considered generics having BP NPs as only kind-referring (e.g., Carlson, 1977; Liebesman, 2011).

(20a) is upward-entailing because it concerns some specific individuals. Some specific tigers can be described, for example, as individual felines, mammals, or animals. The substitution of the term "tiger" in (20a) with the term "animal" in (20b) does not involve a change of the subject referent; thus, (20b) is still a claim about the same specific individuals, this time identified by a different feature. (21a), instead, is not upward-entailing because it concerns a whole kind, not its members. This time, the replacement of the term "tiger" in (21a) with the term "animal" in (21b) involves a change of the subject referent: the species *Panthera tigris* cannot correspond to a more inclusive set – in this case, the biological kingdom *Animalia*. Thus, (21b) concerns a different kind.

It is important to bear in mind that only true and non-negated particular sentences can be used for this test. Indeed, true and non-negated characterizing sentences tend not to be upward-entailing, regardless of their NPs; thus, they cannot serve to distinguish kind-referring from object-referring NPs. As mentioned in the previous subsection, in characterizing sentences genericity is a feature of the whole sentence; in particular, characterizing sentences express generalizations about kinds or their members that might not hold for more inclusive categories. Recall sentence (13b):

(13b) A tiger is striped.

(13b) is a characterizing sentence with an object-referring NP. Such sentence is both true and non-negated; however, it is not upward-entailing. Indeed, substituting its subject term "tiger" with the more inclusive term "animal", as in (22), does not preserve its truth-value:

(22) An animal is striped.

In the next subsection, I consider the functioning of characterizing sentences, by comparing it with the functioning of particular sentences.

1.1.4 Characterizing and particular sentences

Characterizing sentences express generalizations based on properties of kinds, individual members of kinds, or over patterns of events in which the subjects are involved. They are opposed to particular sentences, which concern specific properties of kinds, individuals, or particular episodes. Specifically, Krifka et al. (1995) call the characterizing sentences about properties of kinds or properties of individual members of kinds *lexical characterizing sentences* and the characterizing sentences about patterns of events *habitual sentences*. (23a)-(23c) are examples of lexical characterizing sentences, whereas (23d)-(23f) are examples of habitual sentences:

(23)

a. Chimpanzees are intelligent.

b. The chimpanzee has long arms.

c. A chimpanzee lacks an external tail.

d. Chimpanzees climb trees.

e. A chimpanzee climbs trees.

f. These chimpanzees throw rocks.

The following, instead, are examples of particular sentences:

(24)

a. Primates arose about 85-55 million years ago.

b. Those chimpanzees are throwing rocks.

c. The chimpanzee is outside.

The sentences in (23) and (24) concern kinds or individuals. However, the sentences in (23) express generalizations, whereas the sentences in (24) describe specific facts or episodes.

Importantly, as noted in § 1.1.2, characterizing sentences and kind-referring sentences can overlap. More precisely, characterizing sentences can have kind-referring NPs. While the predicates occurring in characterizing sentences are always stative, kind-referring NPs can be paired with both stative and dynamic predicates. As a consequence, kind-referring NPs can occur in both characterizing and particular sentences. For example, both (23a) and (24a) have a kind-referring NP; however, (23a) expresses a generalization, whereas (24a) concerns a specific fact. Conversely, according to Krifka et al. (1995), characterizing sentences are generics in virtue of the whole sentence and can also contain object-referring NPs. For example, (23c) expresses a generalization over the property of individual chimpanzees.

In the linguistic literature, a central feature attributed to characterizing sentences is *temporal unboundedness* – that is, the truth of characterizing sentences is not bound to specific intervals of times; rather, their truth lacks clear temporal boundaries (see Declerck 1988; Krifka et al., 1995; Mari, Beyssade, & Del Prete, 2013). Indeed, if a characterizing sentence is modified by an adverb denoting specific temporal locations, it becomes odd, as in (25):

(25)

a. Leopards have spots

b. ? Leopards have spots today.

On the contrary, particular sentences can be bounded to specific intervals of time, as shown by sentences (24a) and (26):

(26) The leopard is outside today.

As for kind-referring and object-referring NPs, some linguistic tests have been proposed to disambiguate between characterizing and particular sentences. In the remainder of this subsection, I discuss two tests originally proposed by Krifka et al. (1995).

As a first test, Krifka et al. (1995) proposes to combine the target sentence with an adverb of quantification like *usually*. If modifying the sentence elicits a drastic change in meaning, then the target sentence is a particular sentence; otherwise, it is a characterizing sentence. Krifka et al. (1995) offer the following examples:

(27)

a. A lion has a bushy tail.

b. A lion usually has a bushy tail.

c. A lion stood in front of my tent.

d. A lion usually stood in front of my tent.

When a characterizing sentence like (27a) is combined with the adverb of quantification *usually*, as in (27b), the sentence still expresses a generalization over the tails of lions, while conveying that there may be exceptions. On the contrary, when a particular sentence like (27c) is combined with the same adverb of quantification, as in (27d), the change in meaning is more dramatic; indeed, while (27c) describes an isolated episode, (27d) report a regularity of events.

Let us consider now a second test proposed by Krifka et al. (1995) to disambiguate between characterizing and particular sentences. As they noted, several accounts of generics (e.g., Dahl 1975; Lawler, 1973; Burton-Roberts, 1977; see Cohen, 2001 and Greenberg, 2003 for more recent accounts) agree with the idea that characterizing sentences having an IS NP can only express rules, definitions, norms, central properties, or necessary properties. On the contrary, characterizing sentences containing BP NPs can also express generalizations over accidental properties. To illustrate this contrast, Lawler (1973) offers the following examples:

(28)

a. Madrigals are polyphonic.

b. Madrigals are popular.

c. A madrigal is polyphonic.

d. ? A madrigal is popular.

Being polyphonic is a definitional property for madrigals; for this reason, both (28a) and (28c) are acceptable. However, being popular can only be an accidental property for madrigals. As a consequence, (28b) is acceptable, but (28d) is not. Conversely, consider the examples proposed by Nunberg and Pan (1975):

(29)

a. Football heroes are popular.

b. A football hero is popular.

As Nunberg and Pan (1975) argue, being popular is a central feature for a football hero; thus, both sentences in (29) are acceptable. However, it should be noted that Krifka et al. (1995) do not explicitly state how these observations can be used to distinguish characterizing and particular sentences. A plausible hypothesis is that the purpose of this test is to show that IS sentences about accidental properties are fine only when particular; indeed, a particular sentence like (30) is acceptable:

(30) A madrigal, namely The Silver Swan, is popular.

However, even following this hypothesis, a further complication arises. Contrary to what is assumed in the literature, Krifka (2013) notes that, although IS characterizing sentences tend to be definitional, they can also express descriptive (i.e., non-definitional) generalizations. He offers the following examples:

(31)

a. A trout can be caught by many different methods.

b. A hedgehog makes a good pet.

c. A poodle should be clipped by a professional groomer

d. A madrigal sounds best with all the voice-parts doubled.

The sentences in (31) show that characterizing sentences with an IS NP can have a descriptive interpretation. In particular, some of these statements seem to ascribe accidental properties. For example, "making a good pet" is not necessary for being a hedgehog; nevertheless, (31b) is acceptable. Thus, the problem of how IS characterizing sentences work seems to require further research.

As I cleared from the beginning of this section, my research focuses on the interpretation of *BP characterizing sentences*. In the next subsection, I discuss the reasons for testing this type of generics.

1.1.5 The focus of my research

The category of generics includes various sentences that are similar in their surface form. To avoid potential ambiguities between BP characterizing sentences and other types of sentences, I presented a classification of generics based on the work of Krifka et al. (1995) and other scholars. In this subsection, I discuss why my research focuses specifically on the interpretation of characterizing sentences with BP NPs.

The reason why I focus on characterizing sentences is quite straightforward. I aim at investigating whether the finding of the inferential asymmetry obtained with generic generalizations about animal categories replicates with generic generalizations about social categories. Hence, I am concerned with generalizations. Characterizing sentences express generalizations, therefore; they fall within my area of interest. Sentences with kind-referring NPs, instead, do not always express generalizations, as they are sometimes particular sentences, that concern specific facts and episodes. For this reason, my focus is on characterizing sentences. Characterizing sentences have also attracted most of the attention in the linguistic, philosophical, and psychological literature on generics inasmuch they present puzzling semantic features (see § 1.2) and distinctive cognitive implications (see Chapter 2).

The reasons why my focus is specifically on *BP* characterizing sentences needs a more detailed explanation. To begin with, BP characterizing sentences have a wider distribution compared to both IS and DS characterizing sentences. In other words, they are not affected by the linguistic constraints that, instead, limit IS and DS characterizing sentences. In the previous subsections, I reported two linguistic constraints regarding IS characterizing sentences: their NPs cannot be combined with kind-level predicates unless they receive a taxonomic reading, and they often sound odd when they do not ascribe necessary properties. In particular, the latter constraint is problematic for my research. In contrast to animal kinds, whose members are often highly similar to one another (e.g., different skunks are highly similar in appearance and behavior), members of social categories are more variable (e.g., Brandone, 2017; Nisbett, Krantz, Jepson, & Kunda, 1983). Thus, the properties attributed to members of social categories might be perceived as more accidental than the properties attributed to members of animal kinds. As a consequence, IS social generics might be perceived as odder than IS generics about animals (see also Radden, 2009 on this constraint);[7] in this case, it would be difficult to determine whether any difference between the interpretation of social generics and the interpretation of generics about animals is due to domain or to the IS syntactic form. A further linguistic constraint for IS characterizing sentences has been identified by Greenberg (2003): they can hardly be interpreted as expressing generalizations over the members of a kind when they concern members of *extremely unnatural classes*. The following is one of her examples:

(32) A relatively small lion called "Leo" eats cornflakes for breakfast.

As Greenberg notes, it is easier to interpret IS sentences like (32) as concerning a specific individual rather than the members of a kind. Although she does not focus her analysis on the DS syntactic form, it is possible to extend her observation on this type of sentences. Indeed, (33) seems readily interpretable as concerning a particular individual rather than as expressing a generalization over the members of a kind:

[7] As noted by Radden (2009), IS social generics do not sound odd when they refer to well-defined groups, as in "A customer is always right". However, in my research, I test novel social categories, which cannot be considered as well-defined.

(33) The relatively small lion called “Leo” eats cornflakes for breakfast.

Beyond this shared restriction, DS sentences present a specific nominal constraint, which has been observed by several scholars (e.g., Carlson, 1977; Christophersen, 1939; Krifka et al. 1995; Nunberg & Pan, 1975; Vendler, 1967). Consider the following contrast, originally discussed by Carlson (1977; attributed to Partee):

(34)

a. The coke bottle has a narrow neck.

b. ? The green bottle has a narrow neck.

While (34a) is a regular sentence with a clear meaning, (34b) is odd under a generic reading. Krifka et al. (1995) hypothesized that the nominal constituent of sentences having a DS NP need to refer to a *well-established kind* to be interpreted as kind-referring (the notion of “well-established kind” is not analyzed further in Krifka et al., 1995). Since the nominal constituent of (34b) does not refer to a well-established kind, the sentence sounds odd under a generic reading. This hypothesis is supported by the observation that the corresponding IS sentence is acceptable (although intuitively false):

(35) A green bottle has a narrow neck.

According to Krifka et al. (1995), (35) does not sound odd because, as noted in § 1.1.3, IS NPs on a non-taxonomic reading can only be object-referring; hence, they are acceptable even if they do not concern a well-established kind precisely because they do not refer to kinds. Interestingly, as noted by Radden (2009), this feature of DS sentences explains why, in a generic reading, DS sentences tend to sound odd when used to describe social categories. Here is one of his examples:

(36) ? The Italian loves pasta.

According to Radden (2009), (36) is problematic as a possible generic construction because social categories are usually too heterogeneous to form well-established kinds. This constraint is particularly problematic for my research, as it could add an element of complexity to the items being tested; indeed, in my work I do not test well-established categories. Rather, I use novel categories and novel labels: the purpose of this choice is to ensure that participants do not simply retrieve learned facts. Beyond these linguistic constraints, DS sentences are also the less common type of characterizing sentences, as shown, e.g., by their scarcity in child-directed speech (Gelman, Coley, Rosengren, Hartman, Pappas, & Keil, 1998).

Importantly, BP characterizing sentences are not subject to any of the restrictions that influence IS and DS characterizing sentences; thus, they make a great candidate for my research. First, characterizing sentences with non-taxonomic BP NPs paired with kind-level predicates are well formed, as (37):

(37) Brain-eating amoebas are rare.

Second, BP characterizing sentences ascribing non-necessary properties do not sound odd, as (38):

(38) Hedgehogs make good pets.

Third, as Greenberg (2003) notes, BP characterizing sentences about unnatural classes are still easily interpreted as generalizations over members of a kind, as her BP version of (32) shows:

(39) Relatively small lions called “Leo” eat cornflakes for breakfast.

Finally, as noted by Krifka et al. (1995), contrary to sentences with a DS NP, BP sentences do not need to concern a well-established kind to be interpreted generically. Consider the BP version of (34b):

(40) Green bottles have a narrow neck.

Just like (35), (40) is acceptable. Krifka et al. (1995) conclude from this observation that BP NPs must be ambiguous between two interpretations: they are interpreted as kind-referring when their nominal constituent concerns a well-established kind and as object-referring when their nominal constituent concerns a non-established kind. Since BP characterizing sentences do not need to concern well-established kinds, they do not sound odd when used to describe social categories (though, as I discuss in Chapter 2, they might be problematic for other reasons; see § 2.2). As Radden (2009) observes, while (36) is odd, the following sentence is linguistically fine:

(41) Italians love pasta.

Indeed, (41) expresses a generalization over individual members of a social category; thus, it does not require its nominal constituent to denote a well-established kind.

On top of being free from the linguistic constraints of IS and DS characterizing sentences, I focus on BP characterizing sentences for two other reasons. On the one hand, the previous works on the inferential asymmetry (i.e., Cimpian et al., 2010; Brandone et al., 2015) focused on the interpretation of BP characterizing sentences as well. Thus, investigating the robustness and generalizability of the inferential asymmetry for social generics requires to test the same type of sentences. On the other hand, BP characterizing sentences are outstandingly common in everyday speech: we use them repeatedly, and we hear and read them in newspapers, advertising, movies, social media, songs, and books. They are also markedly frequent in parent-child speech (Gelman, 2003; Gelman et al., 1998; Gelman, Taylor, & Nguyen, 2004).

For all of these reasons, the focus of my dissertation is on BP characterizing sentences. From now on, I will use the term "generics" to refer to this type of sentences exclusively, unless otherwise indicated. In the remainder of this chapter, I first consider the distinction between generics and other non-generic constructions. Then, I discuss two analyses that have been proposed to account for the meaning of generics.

1.2 Generics versus non-generics

Generic sentences express general claims and lack explicit determiners. However, taken singularly, such features are not unique to generics. Sentences including a *quantifier* (e.g., "All ravens are black") can also express generalizations, whereas *existential* sentences (e.g., "Ravens are flying") sometimes lack explicit determiners. In the present section, I discuss the distinction between generics and these non-generic constructions. Such distinction will come in handy for the next section, where I present the problems related to the semantic analysis of generics. The distinction between generics and quantified sentences outlined here will also be useful in Chapter 2, where I discuss the experimental evidence regarding the strength of the two types of generalizations. I first provide a comparison between generics and quantified sentences based on two features possessed by the latter but not by the former. Then, I provide a comparison between generics and existentials, together with some means to distinguish them.

1.2.1 Generics versus quantified sentences

In natural language, generalizations can be expressed either through generic sentences or quantified sentences. Contrary to generics, in quantified sentences, a quantifier – that is, a determiner like "some", "many", "most", or "all" – explicitly states which proportion of Ks, namely the members of the category the sentence is about, have F, namely the predicated property. (41) offers some examples of quantified sentences:

(42)

a. All leopards have spots.

b. Most tigers are striped.

c. Some lions have manes.

More generally, the syntactic forms of the sentences in (42) can be reported as in (43):

(43)

a. All Ks are/have F.

b. Most Ks are/have F.

c. Some Ks are/have F.

The syntactic difference between generics and quantified sentences determines different general linguistic properties for the two types of sentences (this difference also determines different cognitive implications, which I present in Chapter 2). In what follows, I describe two ways in which generics linguistically differ from quantified sentences – that is, a) the former do not have semantic rules simply based on quantity and b) the scope of their nominal argument cannot be contextually restricted in the same way of the latter's.

1.2.1.1 Generics do not simply track quantity

Quantified sentences have precise semantic rules in natural language, which are based on quantity. A sentence like (42a) is true just in case every member of the kind *leopard* has the predicated property. The truth of a sentence like (42b) requires the majority of tigers to be striped. Finally, the sentence (42c) is true if at least one lion has a mane. In formal semantics, these rules are represented as relations between sets (e.g., Barwise & Cooper, 1981). For example, the quantifier "all" expresses an inclusion relation: (42a) is true if and only if the set of leopards is a subset of the set of things that have spots. Conversely, the quantifier "some" expresses an overlap relation: for the sentence (42c) to be true, the intersection between the set of lions and the set of things that have manes must be non-null.

At first glance, generics seem to present semantic rules similar to the ones concerning quantified sentences. For example, consider the following pair (Carlson, 1977):

(44)

a. All dogs are mammals.

b. Dogs are mammals.

(44a) and (44b) seem to have similar truth-conditions; indeed, both (44a) and (44b) are true in the same circumstances (that is, they are both true just in case every dog is a mammal). However, a closer analysis reveals that the truth-conditions of most generics do not conform to this case. In contrast to quantified sentences, generics do not simply track quantity (e.g., Abelson & Kanouse, 1966; Brandone et al., 2015; Carlson, 1977; Cimpian et al., 2010; Gelman, Star, & Flukes, 2002; Krifka et al., 1995; Leslie, 2008). An intriguing feature of generics is that they often gloss over exceptions. For example, consider the following pair:

(45)

a. All ravens are black.

b. Ravens are black.

Universally quantified sentence (45a) is false; indeed, not all existing ravens are black, as albino ravens exist in nature. However, generic (45b) is still intuitively true, despite the existence of the same non-conforming category members. This comparison shows that universally quantified sentences and generics do not have similar truth-conditions. One could then hypothesize that the latter simply have slightly more flexible truth-conditions than the former. In particular, for a generic to be true, it may be sufficient that the majority of Ks have F. For example, consider the following most-quantified sentence:

(46) Most ravens are black.

In contrast to (45a), (46) is true; indeed, the truth of sentences including the quantifier "most" can tolerate a number of exceptions, as long as the non-conforming Ks are less than half. Generic (45b) and most-quantified sentence (46) are true in the same circumstances; thus, generics and most-quantified sentences might have similar truth-conditions. However, this hypothesis does not hold; indeed, other generic sentences are intuitively true despite the proportion of Ks having F is less than 50%. For example, consider the following pair:

(47)

a. Most ducks lay eggs.

b. Ducks lay eggs.

(47a) is false; indeed, only a small subset of ducks actually lay eggs (that is, the female fertile ducks). In contrast to (47a), generic (47b) is still intuitively true; therefore, the hypothesis that generics and most-quantified sentences have similar truth-conditions should be rejected. At this point, one might adopt an even weaker hypothesis, namely that generics' truth-conditions are as flexible as the ones of some-quantified sentences. The truth of some-quantified sentences requires that at least one K have F; thus, they are true even when most (but not all) of Ks are not conforming. As generic (47b), the following some-quantified sentence is true:

(48) Some ducks lay eggs.

However, consider now the following pair:

(49)

a. Some ducks are females.

b. Ducks are females.

Some-quantified sentence (49a) is true; since a great proportion of ducks are female, the condition that at least one K needs to have F is satisfied. Generic (49b), instead, is intuitively

false. This observation is particularly compelling since, conversely, generic (47b) is intuitively true, despite the number of female ducks is greater than that of egg-laying ducks. Therefore, as these comparisons highlight, the hypothesis that generics and some-quantified sentences have similar truth-conditions must be rejected. That some Ks have F does not constitute a sufficient condition to account for generics' truth and falsity. Further evidence for this is provided by the comparison between the following sentences, originally proposed by Leslie (2007, 2008):

(50)

a. Mosquitoes carry the West Nile virus.

b. Mosquitoes don't carry the West Nile virus.

As Leslie notices, generic (50a) is intuitively true, despite only 1% of mosquitoes carry the West Nile virus, and only in affected areas (Cox, 2004). Conversely, generic (50b) is intuitively false, even though more than 99% of Mosquitoes do not have the ascribed property. Consider another set of sentences, this time about humans:

(51)

a. Most humans are right-handed.

b. Some humans are left-handed.

c. Humans are right-handed.

d. Humans are left-handed.

Both the quantified sentences (51a) and (51b) are true; indeed, the majority of humans are right-handed, whereas some of them are left-handed. Nevertheless, both generics (51c) and (51d) are intuitively false. Surprisingly, their negation is also intuitively false, as for the negation of the generic (49b):

(52)

a. Humans aren't right-handed.

b. Humans aren't left-handed.

c. Ducks aren't female.

In conclusion, the examples discussed highlight how generics' truth-conditions, contrary to quantified sentences', are not simply based on the proportion of Ks displaying F. This was also well-illustrated by Carlson (1977), who notices that generics cannot be used to answers to questions such as "How many Ks are F?". To use his examples, consider the following question:

(53) How many people eat squid?

Any of the following quantified sentences can answer question (53):

(54)

a. All people eat squid.

b. Most people eat squid

c. Some people eat squid.

Regardless of their truth-value, all the sentences in (54) are appropriate answers to question (53). However, this question cannot be answered by a generic. For example, the following sentence is not an appropriate answer to the question (53):

(55) People eat squid.

1.2.1.2 Generics resist contextual restriction

Depending on the conversational context, the nominal argument of a quantified sentence can be subject to the phenomenon known as *quantifier domain restriction* (see Stalnaker, 1970;

Stanley & Szabó, 2000). Quantifier domain restriction occurs when the scope of the nominal argument of a quantified sentence is restricted to a proportion of Ks made salient by the context of utterance. For example, consider the following sentence – uttered by an ethologist while observing a group of hyenas:

(56) All hyenas are quiet.

In this context, the utterance is not conveying something concerning every existing hyena; rather, it is conveying something concerning the specific hyenas observed by the speaker. The mere presence of salient hyenas in the conversational context is sufficient to make this restricted interpretation available. In contrast to the interpretation of quantified sentences', generics' nominal argument resists contextual restriction (e.g., Dahl, 1975; Krifka, 1987; Krifka et al., 1995; Lazaridou-Chatzigoga, Katsos, & Stockall, 2015; Nickel, 2016). For example, consider the following generic:

(57) Hyenas are quiet.

Uttered in the same context, generic (57) cannot be interpreted as referring solely to the present hyenas; rather, it is still conveying something about hyenas in general. In this case, the presence of the same salient Ks is not sufficient to make a restricted interpretation of the utterance available. Of course, as noted by several scholars (e.g., Cavedon & Glasbey, 1996; Greenberg, 2003; Nickel, 2008, 2016; Haslanger, 2014; Sterken, 2015), in certain circumstances context can play an important role in the interpretation of generics. To illustrate this, Nickel (2008, 2016) proposes to compare the interpretation of the following generics in different contexts:

(58)

a. Dobermans have floppy ears.

b. Dobermans have pointy ears.

(58a) is intuitively true when uttered in a context where the biology of dogs is being discussed; in the same context, (58b) is intuitively false. Conversely, (58a) is intuitively false when uttered in a conversation about dog breeding; in this context, however, (58b) is intuitively true.[8] Although generics can be context-sensitive, they are not in the same way as quantified sentences. Indeed, as illustrated by the interpretation of (57), the presence of salient Ks in the conversational context is not sufficient to make a restricted interpretation of generics' nominal argument available.

1.2.2 Generics versus existentials

Not every sentence having a BP NP is a generic. Existential sentences, or existentials, can have the same syntactic NP; yet, contrary to generics, they express propositions about the existence or presence of some specific individuals. Consider the following sentences:

(59)

a. Grasshoppers are eating the harvest.

b. Grasshoppers are herbivorous.

(59a) is an example of an existential, whereas (59b) is a generic. Contrary to (59b), which expresses a generalization over grasshoppers, (59a) makes no generalizations; rather, it concerns some specific grasshoppers. Sentences like (59a) are called existentials because their nominal argument can be paraphrased using an *existential quantifier* – that is, a quantifier like "some" or "there are" – without changing their meaning. For example, (59a) can be paraphrased as follows:

(60)

a. Some grasshoppers are eating the harvest.

b. There are grasshoppers eating the harvest.

[8] It should be noted, though, that this example works only in contexts where ear cropping is still considered legal. Fortunately, the practice has been banned in several countries.

(59a), (60a), and (60b) have the same truth-conditions; indeed, they are true in the same circumstances. As discussed above (see § 1.2.1.1), generics cannot be paraphrased in the same way; indeed, that some Ks have F does not constitute a sufficient condition to account for their truth or falsity.

Contrary to the distinction between generics and quantified sentences, it is not possible to distinguish a generic from an existential simply based on the syntactic form of their NPs. One has to think about the context and the intended referent as well. In order not to mistake non-generic sentences for, e.g., possible counterexamples to a semantic account of generics, it is always important to verify that we are considering genuine generics, rather than sentences that simply have a similar syntactic form. In this regard, some linguistic tools can help.

A first criterion to distinguish a generic from an existential is provided by Carlson (1977), who proposes to focus on the predicates occurring in these sentences. According to Carlson, the generic or the existential reading of a sentence is determined by whether its NP is paired with an *individual level predicate* or a *stage level predicate*. Individual level predicates, or *i-predicates*, tend to ascribe properties that are stable over times, whereas stage level predicates, or s-*predicates*, tend to ascribe temporary bounded states. Carlson (1977) proposes this distinction based on an ontological criterion. In his view, properties (e.g., "Ks run") are predicated of individuals, whereas states (e.g., "Ks are running") are predicated of stages of individuals – that is, temporary bounded proportion of those individuals. Examples of i-predicates are "run", "know Japanese", "are herbivorous", "are intelligent", "are altruistic". Conversely, "are running", "are speaking Japanese", "are eating", "are hungry", and "are nervous" are examples of s-predicates. Carlson proposes that i-predicates favor the generic reading of BP sentences, whereas s-predicates favor the existential reading of the same sentences. Concerning the sentences in (59), the progressive "are eating" in (59a) expresses a state of grasshoppers; for this reason, the sentence receives an existential reading. The predicate "are herbivorous" in (59b), instead, express a property that is stable over time;

thus, (59b) receives a generic reading.[9] There are several linguistic tests that can be used to differentiate between i-predicates and s-predicates (e.g., Carlson, 1977; Chierchia, 1995; Fernald, 2000; Kratzer, 1995; Milsark, 1974; Stump, 1985). For example, s-predicates can be combined with locative modifiers; however, i-predicates cannot. Thus, pairing a predicate with a locative modifier can diagnoses whether it is an i- or a s- predicate. Consider the following sentences:

(61)

a. Grasshoppers are eating the harvest in the barn.

b. # Grasshoppers are herbivorous in the barn.

Contrary to (61a), (61b) is ill-formed. As a consequence, the predicate occurring in (61a) can be classified as a s-predicate; the predicate occurring in (61b), instead, can be classified as an i-predicate.

Other criteria to distinguish generics from existentials are provided by Leslie (2012a) and Leslie and Lerner (2016). In their view, two well-known linguistic tests can be used to distinguish the two kinds of sentences – that is, the upward-entailment test described in § 1.1.3 (e.g., Krifka et al. 1995; Laca, 1990; Lawler, 1973) and the insertion of the adverb of quantification test presented in §1.1.4 (Krifka et al., 1995). Let us begin with the former.

According to Leslie (2012a) and Leslie & Lerner (2016), existentials are upward-entailing; assuming their truth, they remain true even if more inclusive categories replace their subject categories. For example, the following sentence shows that the term "grasshoppers" in (59a) can be substituted by the term "animals" salva veritate:

[9] The use of the distinction between i- and s- predicates to differentiate between generics and existentials is not without controversy. For example, as highlighted by Leslie & Lerner (2016), such distinction cannot account for the interpretation of DS sentences; depending on the context, to use one of their examples, the sentence "The tiger is intelligent" can be both interpreted generically or existentially. In certain conversational contexts, a certain tiger could be so salient as to inhibit the generic reading of this sentence, thus making the existential interpretation available; however, "is intelligent" is an i-predicate. This distinction also cannot account for the interpretation of IS sentences; as previously discussed (see § 1.1.5), IS sentences can hardly be interpreted as generics when they describe Ks belonging to extremely unnatural classes (Greenberg, 2003), even if their NPs are paired with i-predicates. Given the focus of this work on BP generics, these issues can be set aside. Another interesting problem concerns the reasons why i- and s- predicates give rise to generic and existential readings, respectively. The problem is still up for debate; however, its analysis is beyond the scope of the present work (for more discussion, see Chierchia, 1995; Cohen & Erteschik-Shir, 2002; Kratzer, 1995).

(62) Animals are eating the harvest.

In contrast, this does not apply to generics; indeed, as noted in § 1.1.3, they tend not to be upward-entailing. For example, while the generic (59b) is true, the following sentence is false:

(63) Animals are herbivorous.

Thus, Leslie (2012a) and Leslie and Lerner (2016) argue that verifying whether a sentence is upward-entailing can be used as a criterion to distinguish generics from existentials. However, as noted by Krifka et al. (1995), only particular sentences can be used for this test, with the aim to distinguish kind-referring from object-referring NPs. To be used for the purpose intended by Leslie (2012a) and Leslie and Lerner (2016), this test may require further conditions. Indeed, while generics tend not to be upward-entailing, they sometimes are. For example, consider the following sentence:

(64) Dogs are animals.

This generic is upward-entailing; indeed, if we replace its subject category with the same category used in the previous example, we obtain another true generic:

(65) Animals are animals.

The second test that Leslie (2012a) and Leslie and Lerner (2016) consider has been proposed by Krifka et al. (1995) to distinguish between characterizing and particular sentences. According to this test (see § 1.1.4), if a sentence modified with an adverb of quantification exhibits a drastic change in meaning, the examined sentence is particular – otherwise, it is characterizing. Leslie (2012a) and Leslie & Lerner (2016) assume that this test can be also used to distinguish generics from existentials. For example, consider the following sentences:

(66)

a. Grasshoppers are usually eating the harvest.

b. Grasshoppers are usually herbivorous.

The sentences in (66a) and (66b) are a modified version of the existential (59a) and the generic (59b), respectively. While the existential combined with the adverb of quantification express now a regularity of events, the generic modified in the same way still expresses a generalization.

1.3 What do generics mean?

Generics are frequent in natural language and play a powerful role in children's development of kind concepts (e.g., Brandone, Cimpian, Leslie, & Gelman, 2012; Gelman, 2003; Gelman & Bloom, 2007; Rhodes, Leslie, Bianchi, & Chalik, 2018a; Segall, Birnbaum, Deeb, & Diesendruck, 2015; Sutherland & Cimpian, 2015). Nonetheless, the semantics of generics is not straightforward. As discussed in the previous section, both generics and quantified sentences express generalizations; yet, in contrast to quantified sentences, generics are not simply about quantity and resist contextual restriction. Likewise, both generics and some existentials lack explicit determiners; however, in contrast to existentials, generics do not express propositions about specific individuals. While the truth-conditions of quantified sentences and existentials are straightforward, determining those of generics has proven challenging. Starting from the earlier attempts of the seventies (e.g., Carlson, 1977; Dahl, 1975; Lawler, 1973), a great deal of research has been conducted on the semantics of generics; yet, their meaning is still debated nowadays.

Although its assumptions have recently been criticized (e.g., Liebesman, 2011), the standard view in the debate is that the logical form of generics contains an operator contributing to their compositional meaning but covert from the surface form (e.g., Heim, 1982; Krifka et al. 1995; Greenberg, 2003). Therefore, the standard theories try to account for generics' meaning by explaining the functioning of this hidden operator. More recently,

other less formal approaches started to give more emphasis on the role that human cognition plays in the interpretation of generics (e.g., Leslie, 2008; Prasada, 2000). In these views, generics function as a linguistic outlet of how we conceptualize the relations between categories and properties.

In the first part of this section, I provide a brief overview of some of the semantic theories that have been proposed to account for generics' meaning. In the second part, I present a formal analysis of generics – that is, the probabilistic account put forward by Cohen (e.g., 1996, 1999, 2004). In the third part, instead, I present a cognitive account of generics – that is, the content-based account proposed by Leslie (e.g., 2007, 2008, 2015).

1.3.1 Semantic analyses of generics: an overview

Despite generics NPs' lack of explicit quantifiers, most formal semantic treatments of generics try to account for their meaning by positing a covert two-place operator in their nominal argument, called *GEN*. According to this standard view, GEN is not phonologically realized; nonetheless, it would belong to the logical form of generics and contribute to their compositional meaning. Specifically, as first proposed by Heim (1982), Farkas (1981), Farkas and Sugioka (1983), and Schubert and Pelletier (1989), generics' logical form is analyzed as tripartite structure consisting of the operator GEN, a restrictor R, and a scope S. The restrictor R's role is to pick out the domain of GEN; the scope S, instead, picks out the properties attributed to GEN elements of R. The whole structure can be represented as follows:

GEN [R] [S] [10]

[10] Such structure relies on Lewis' (1975) account of the logical form of constructions containing explicit adverbs of quantification. According to Lewis, adverbs of quantifications are unselective: they bind any free variable in the sentence in their scope. In describing the possible logical form of generics, a similar operator was introduced to overcome some difficulties presented by Carlson's (1977) first proposal of a monadic type-shifting operator *Gn* (for more discussion, see Krifka et al., 1995: 23–30).

Accordingly, the logical form of a generic like "Killer whales are intelligent" corresponds to the following:

GEN(x) [Killer whale(x)] [Intelligent(x)]

Therefore, to model the truth-conditions of generics, the strategy of the semantic theories positing GEN is to provide an account of its functioning. For the reasons discussed in § 1.2.1, such accounts cannot simply analyze GEN in terms of one of the standard quantifiers. Thus, they need to characterize this operator in alternative ways, which could account for our intuitions on the truth and falsity of such sentences. These approaches can be divided into several families, depending on the type of characterization they provide of GEN. Below, I provide a brief list of some of the most well-known families, which is not exhaustive.

Many accounts (e.g., Chierchia, 1995; Greenberg, 2003; Heim, 1982; Krifka et al. 1995; Schubert and Pelletier, 1989) treat GEN as modalized operator quantifying over possible worlds. Specifically, some authors characterize such possible worlds as worlds where things go as normally as possible for a given K (e.g., Asher & Morreau, 1995; Pelletier and Asher, 1997). Roughly speaking, the general idea behind these accounts is that a generic like "Ravens are black" is true despite the existence of non-conforming individuals because, in a world where things go as normally as possible, every raven would have the ascribed property. Similarly, and more recently, Nickel (2008, 2016) analyzes GEN as an operator quantifying over normal Ks; in his view, "Ravens are black" is true if and only if «there is a way of being a normally colored raven such that all ravens normal in that way are black» (Nickel, 2016: 63). Other accounts argue for the importance of the context on the interpretation of generics; for example, Declerck (1991) proposes that GEN quantifies over contextually-relevant individuals, whereas Sterken (2015) proposes that GEN functions as an indexical constituted by a quantificational force and a lexical restriction, both determined by the context.[11] Other authors argues that GEN expresses stereotypical (e.g., Declerck, 1986; Geurts, 1985) or

[11] Indexicals are expressions like "I" and "now", whose meaning can change depending on the speaker who use them and the context of utterance (see Kaplan, 1989).

prototypical knowledge (e.g., Heyer, 1985; Nunberg and Pan, 1975; Platteau, 1980).[12] Finally, some authors contend that GEN concerns probability judgments (e.g., Cohen, 1996, 1999, 2004; Tessler & Goodman, 2019). For example, roughly speaking, Cohen proposes that generics are true just in case the probability that an arbitrary K has F is higher than two possible thresholds (see the next subsection).

The formal accounts deviating from the standard view argue against the possibility that generics' logical form contains a covert, unpronounced operator. For example, according to Liebesman (2011), generics are simple subject/predicate sentences ascribing properties to kinds. Other less formal approaches (e.g., Leslie, 2008; Prasada, 2000) argue that to investigate the meaning of generics we should consider the role that such sentences play in our cognition. Such cognitive approaches share the assumption that the use and interpretation of generics have a peculiar way to influence and interact with our conceptual knowledge – that is, our ability to develop, recall, and grasp concepts and classifications about the world. In particular, according to these views, the flexibility of generics' truth-conditions indicates that these sentences express several kinds of connections between categories and properties, depending on the type of property being ascribed.

The theories mentioned in this subsection have the merit to capture several intuitions that we have on generics' meaning. Of course, they are also far more complex and detailed than reported here. However, despite all the work that has been done on the topic, there is still no theory capable of fully accounting for the odd semantic functioning of generics. In the next subsections, I present two recent and well-known account of generics, one belonging to the standard view and the other to the cognitive approach: the probabilistic account of Cohen (e.g., 1996, 1999, 2004) and the content-based account of Leslie (e.g., 2007, 2008, 2015).

[12] Referring to the notion of "prototype" developed by Rosch (1978) and Rosch and Mervis (1975), the prototype approaches argue that generics concern the members that are the *best* or *more central* representations of a certain category. In other words, these approaches analyze GEN as quantifying over the prototypical members of a category.

1.3.2 A probability-based approach

Cohen (1996, 1999, 2004) proposes a sophisticated formal account of generics based on probabilistic considerations. Herein, I provide an overview of his theory.

Cohen's proposal is that a generic "Ks are F" is true just in case the probability that an arbitrary K has F is higher than two possible thresholds *r*. More precisely, according to Cohen, the arbitrary K must belong to a relevant subset of Ks – that is, the Ks that satisfy at least one property in ALT(F). Cohen defines ALT(F) as the set of properties that are co-alternative to F. For example, consider the following generic:

(67) Tigers are striped.

In Cohen's view, a generic like (67) is evaluated with respect to the tigers that satisfy at least one of the co-alternative properties to "being striped". (67) ascribes a type of fur. In this case, ALT(F) contains the following properties:

ALT(F) = {being striped, being spotted, being uniform, …}

The relevant tigers for the interpretation of (67) are all tigers having a type of fur; thus, (67) is true just in case the probability that an arbitrary tiger with a fur is striped is higher than r. As mentioned above, Cohen considers two possible thresholds r. The suitable threshold is determined by the type of reading that a generic receives.

According to Cohen, generics can have an *independent* reading or a *dependent* reading. Generics receiving an independent reading have a fixed r of 0.5. Therefore, under this reading, generics' truth-conditions are the following (adapted from Cohen, 1996: 39):

"Ks are F" is true just in case the probability that an arbitrary K that satisfies some predicate in ALT(F) has F is greater than 0.5.

Consider again (67). In Cohen's view, a generic like (67) receives an independent reading, and it is correctly predicted to be true. Indeed, the probability that an arbitrary tiger having a fur is striped is greater than 0.5.

For generics receiving a dependent reading, instead, r is based on a comparison with an arbitrary member of ALT(K), which is the set of categories that are co-alternative to K. Specifically, under this second reading, r is the probability that an arbitrary member of ALT(K) that satisfies some predicate in ALT(F) has of satisfying F. Thus, the truth-conditions of a generic receiving a dependent reading are the following (adapted from Cohen, 1996: 60):

"Ks are F" is true just in case the probability that an arbitrary K that satisfies some predicate in ALT(F) has F is greater than the probability that an arbitrary member of ALT(K) that satisfies some predicate in ALT(F) has of satisfying F.

As an example of how such analysis works, consider now the following generic:

(68) Ticks carry the Lyme disease.

(68) concerns a disease transmissible by an insect. This time, ALT(F) contains the following properties:

ALT(F) = {carrying the Lyme disease, carrying malaria, carrying the yellow fever, …}

Accordingly, the relevant Ks are the disease-carrying ticks. ALT(K), instead, contain other disease-carrying insects that are co-alternative to ticks:

ALT(K) = {ticks, mosquitoes, flies, …}

Thus, according to Cohen, a generic like (68) is true just in case the probability that an arbitrary disease-carrying tick carries the Lyme disease is greater than the probability that an arbitrary disease-carrying insect carries the Lyme disease. Since this is the case, then (68) is correctly predicted to be true.

Cohen's proposal has the merit to account for our intuitions about the truth and falsity of a great variety of generics. Consider the following sentences, which have been used to discuss the flexibility of generics' truth-conditions in § 1.2.1.1:

(44b) Dogs are mammals.
(45b) Ravens are black.
(47b) Ducks lay eggs.
(50a) Mosquitoes carry the West Nile virus.

Cohen's account correctly predicts these generics to be true. (44b), (45b), and (47b) receive an independent reading. Concerning (44b), the probability that an arbitrary dog that satisfies some predicate in ALT(F) is a mammal is higher than 0.5. Consider (45b), instead. Despite the existence of albino ravens, the probability that an arbitrary raven that satisfies some predicate in ALT(F) (e.g., "being black" and "being albino") is black is higher than 0.5. Regarding (47b), the probability that an arbitrary gestating duck that satisfies some predicate in ALT(F) (e.g., "laying eggs", "bringing forth live young", and "undergoing mitosis") lays eggs is, again, higher than 0.5. (50a), instead, receives a dependent reading. Similarly to the analysis of (68), an arbitrary disease-carrying mosquito has a probability of carrying the West-Nile virus greater than the probability that an arbitrary disease-carrying insect has to carry the same disease. Hence, (50a) is correctly predicted to be true as well.

However, as noted by Cohen, his proposal, as it stands, is too permissible, and it yields incorrect predictions. Here is one of his examples, which receives an independent reading:

(69) Bees are sexually sterile.

At this stage, Cohen's analysis incorrectly predicts that (69) is true. Indeed, most bees are sterile, and the probability that an arbitrary bee that satisfies some predicate in ALT(F) (i.e., "being sterile" and "being fertile") is sterile is higher than 0.5. Generics like (70) are also problematic:

(70) Bees are female.

In the present condition, since most bees are females, Cohen's account incorrectly predicts that (70) is true as well. To render his analysis less permissible, Cohen set a further condition on his semantic account, which he labels *homogeneity constraint*. According to this additional requirement, the probability conditions illustrated above must hold in all the *salient partitions* of the category K (readapted from Cohen, 1996: 88 – 89).[13] The homogeneity constraint accounts for the falsity of (69); indeed, there is a salient partition of bees in castes where the ability to reproduce is differently distributed: while an arbitrary bee has a probability of being sterile higher than 0.5, a drone or a queen bee does not. The homogeneity constraint also accounts for the falsity of (70): while an arbitrary bee has a probability of being female higher than 0.5, a drone does not. The probability conditions required in evaluating (69) and (70) do not hold for all the salient partition of bees. Due to the addition of the homogeneity constraint, both (69) and (70) are now correctly predicted to be false.

Cohen's theory is one of the analysis with the greatest explanatory power in the debate, as it can account for a wide range of generic sentences. However, it also encounters some difficulties. As noted by Leslie (2008), a first problem with Cohen's proposal is that it does not include a criterion to determine whether a given generic receive an independent or a dependent reading. Leslie (2008) also highlights other issues related to the homogeneity constraint. Cohen introduces this constraint to prevent his semantic account from being too

[13] According to Cohen, what counts as a salient partition in a certain context might vary depending on the culture, the language, and the individuals we are considering. The notion is vague, but it could account for the particular way in which generics can be contextually restricted (see § 1.2.1.2). As he writes, «[f]or a zoologist, very fine distinctions among birds might be salient; not so for the layperson. Indeed, the zoologist might judge [Birds fly] to be false for just this reason» (Cohen, 1996: 91).

permissible. Such modification enables him to correctly predict generics like (69) to be false. However, the homogeneity constraint also makes the following generic false:

(71) Bees reproduce.

Indeed, a salient partition of bees is composed of sterile bees – that is, the workers. Furthermore, according to Leslie (2007, 2008), due to the introduction of the homogeneity constraint, sex-specific generics as (47b) also become problematic:

(47b) Ducks lay eggs.

Indeed, the physiological sex of an animal is a salient partition in the evaluation of sex-specific generics. According to Leslie, both male ducks and female ducks are salient partitions in the evaluation of (47b); however, since only female ducks can lay eggs, then the account incorrectly predicts (47b) to be false. However, this potential counterexample to Cohen's theory can be overcome by noting that sentences like (47b) are evaluated with respect to the relevant Ks – in this case, the gestating ducks. This subset of ducks cannot be partitioned according to sex, nor it can be partitioned between fertile and sterile ducks since only the fertile female ducks can satisfy some predicate in ALT(F) (e.g., "laying eggs", "bringing forth live young", and "undergoing mitosis"). Due to the absence of a related salient partition, (47b) is correctly predicted to be true despite the addition of the homogeneity constraint. Nevertheless, as the case of (71) shows, the addition of the homogeneity constraint might make Cohen's analysis too rigid to account for generics' truth-conditions. Other counterexamples proposed by Leslie (2007, 2008) illustrate how the dependent reading, instead, might be too flexible. These counterexamples are generics concerning impairments that are more frequent among Ks than among the members of a co-alternative kind. Here is one of Leslie's examples:

(72) Humans are autistic.

Under a dependent reading, (72) is incorrectly predicted to be true. Indeed, compared to other species, humans are more likely to display autistic-like traits. However, generics like (72) are intuitively false.

Moreover, it seems to me that similar problems affect the independent reading. Recall, for example, (51c):

(51c) Humans are right-handed.

(51c) is intuitively false, although the majority of humans are right-handed. However, Cohen's account incorrectly predicts (51c) to be true; indeed, the probability that an arbitrary human is right-handed is higher than 0.5, and there appears to be no salient partitions where this condition is not satisfied.

Finally, a further issue with Cohen's proposal is highlighted by Nickel (2008, 2016), who notes that it cannot account for generics ascribing multiple properties that are mutually exclusive. Here is one of his examples:

(73) Elephants live in Africa and Asia.

According to Nickel, a generic like (73) is semantically equivalent to the following sentence:

(74) Elephants live in Africa and elephants live in Asia.

However, it is simply not possible that most elephants live in Africa and that most elephants live in Asia (unless we formulate the unlikely hypothesis that most elephants spend their life partly in Africa and partly in Asia). Thus, the truth of generics like (74) is troublesome for Cohen's account.

1.3.3 A content-based approach

Leslie (e.g., 2007, 2008, 2015) proposes an innovative account of generics that relies on content-based factors. Although her theory encounters some difficulties, it has been extremely important for the recent development of the debate. First, she provides a theory of generics' meaning that overcomes several problems that instead afflict other analyses. Second, she characterizes generics in a way that can be experimentally investigated by the cognitive sciences; in turn, this generated a consistent set of experimental works based on her original proposal. Finally, Leslie's analysis has allowed the flourishing of numerous works in philosophy of language and psychology on the use of generics as stereotypes. In the present section, I address Leslie's treatment of generics' truth-conditions. In the next chapter, instead, I discuss the importance of her account for the cognitive research on generics, along with her hypothesis that the use and interpretation of generics reflect the functioning of a *default mode of generalizing* cognitive mechanism (Leslie, 2007, 2008, 2012b; see also Gelman, 2003, 2010); in the next chapter, I also underlie the importance of her theory for the study of social generics.

According to Leslie (2007, 2008), generics are a linguistic outlet of how we organize our representation of categories. In particular, she claims that generics' variable truth-conditions are determined by the type of property that they predicate. Different types of properties have different informational value and, in turn, are diversely prominent in our conceptual knowledge; thus, in Leslie's view, properties having different informational value determine different truth-conditions for generics. In other words, while quantified sentences' truth-conditions are based on the proportion of Ks displaying F, generics' truth-conditions are based on the content of F. According to Leslie, generics can be divided into four categories depending on the types of properties they predicate: *characteristic*, *striking*, *common*, or *normative* (the last one added more recently; see Leslie, 2015).

Characteristic generics ascribe a characteristic dimension to Ks. According to Leslie, the use and the interpretation of such generics reflect particular background expectations toward certain kinds. For example, we expect the members of an animal category to have a specific mode of reproduction, locomotion, and a distinctive sound. For artifacts and social

categories, instead, we have background expectations concerning their role, purpose, and function. The following are some of her examples:

(75)

a. Bees lay eggs.

b. OrangeCrusher2000s crush oranges.

c. Firemen fight fires.

According to Leslie, we do not rely on statistical frequency to accept these generics; for example, since we have the background expectation that an animal category has a specific mode of reproduction, we only need to observe a few egg-laying bees to accept (75a). In Leslie's view, characteristic generics about animals are true just in case they meet the following conditions: a) some Ks have F, b) F lies along a characteristic dimension for the Ks, and c) the non-conforming Ks are only *negative counterinstances* – that is, the non-conforming Ks fail to have F without possessing an alternative property that is incompatible with F. For example, (75a) is correctly predicted to be true by Leslie's account inasmuch all the three conditions are satisfied. First, some bees lay eggs. Second, since we have the background expectation that an animal category has a specific mode of reproduction, "laying eggs" lies along a characteristic dimension (i.e., "mode of reproduction") for bees. Finally, although almost all bees do not lay eggs, the non-conforming members do not possess an alternative property that is incompatible with "laying eggs" (e.g., "bringing forth live young"); thus, (75a) is accepted and retained true despite a majority of non-conforming members. However, characteristic generics are false if some non-conforming Ks are *positive counterinstances* – that is, in case some non-conforming Ks that fail to have F display an alternative property that is incompatible with F. For example, consider again (70):

(70) Bees are female.

According to Leslie, (70) is false inasmuch the non-conforming bees possess an alternative property that is incompatible with "being female" – that is, "being male". While

male bees are compatible counterinstances for the truth of (75a), they are not compatible for truth of (70). In Leslie's view, the truth-conditions for characteristic generics about artifacts and social categories are similar to the ones of characteristic generics about animals, with the exception that their truth does not require that some Ks have F. This would account for the reason why a generic like (75b) is intuitively true even in the case that no OrangeCrusher2000 has ever been or will ever be used; furthermore, it would also account for the reason why (75c) would be intuitively true even in the case that the climate would be so humid that no fire could start.

Striking property generics ascribe properties that are perceived as particularly appalling and dangerous. (50a) is an example of a striking property generic:

(50a) Mosquitoes carry the West Nile virus.

According to Leslie, such generics are extremely tolerant to counterinstances. For example, as already discussed (see § 1.2.1.1), Leslie notes that (50a) is intuitively true though less than 1% of Ks have the ascribed property (Cox, 2004). Leslie claims that striking property generics are true just in case they meet the following conditions: a) the non-conforming Ks are negative counterinstances, b) some Ks have F, and c) being a K is a *good predictor* of F – that is, the non-conforming Ks need to be disposed to have F under the right circumstances. According to Leslie, a generic like (50a) is true despite the vast majority of non-conforming members because all of these conditions are satisfied. First, no mosquito possesses an alternative property to "carrying the West Nile virus". In Leslie's view, a property like "carrying malaria" does not count as an alternative property; since a mosquito could carry both, the two properties are compatible. Second, some mosquitoes carry the relevant virus. Finally, mosquitoes are good predictors of the ascribed property; indeed, non-conforming mosquitoes are disposed to carry the West Nile virus due to their mobility and blood-feeding behavior. For example, they would carry the relevant virus if they were feeding on infected animals. Consider, instead, the following generic:

(76) Insects carry the West Nile virus.

Leslie's account correctly predicts (76) to be false; indeed, while some insects are mosquitoes, insects do not have a shared disposition to carry the West Nile virus like mosquitoes do. Thus, (76) do not satisfy the relevant condition. More recently, Leslie (2017) noted that we do not tend to accept striking property generics when they describe categories that are socially close to us. This would explain why we intuitively accept generics like (77a) but reject generics like (77b):

(77)

a. Sharks kill people.

b. Dogs kill people.

Dog attacks cause more deaths than shark attacks.[14] However, (77a) is intuitively true, whereas (77b) is intuitively false. According to Leslie, in cases like these, we seek more restricted generalizations involving what we perceive as better predictors. Thus, we might be more likely to accept the following restricted version of (77b):

(78) Poorly trained pit bulls kill people.

Likewise, due to a greater familiarity with sharks, a marine biologist might reject (77a); indeed, most shark species are not dangerous for humans. A marine biologist might instead accept the following restricted versions of (77a):

(79)

a. Bull sharks kill people.

b. Great white sharks kill people.

[14] According to the International Shark Attack File (2014), between 2001 and 2013, in the U.S., dogs killed 364 people, whereas sharks killed 11 people.

Majority generics highlight statistical frequencies and ascribe properties that are less prominent in our conceptual knowledge than other generics. According to Leslie, these generics concern neutral information (neither characteristics nor striking) about Ks and are true just in case a) most Ks have F and b) the non-conforming Ks are negative counterinstances. Here is one of Leslie's examples:

(80) Cars have radios.

Leslie's account correctly predicts (80) to be true; indeed, most cars have radios, and no car has an alternative property that is incompatible with "having radios".

Finally, *normative generics* do not express generalizations about the empirical world; rather, they have a normative force. In other words, instead of describing how Ks are, they describe how Ks *should be* based on cultural norms. Many extremely problematic gender-specific stereotypes fall into this category. The following are some of the examples provided by Leslie (2015):

(81)

a. Boys don't cry.

b. Women put their families before their career.

As highlighted by Leslie, normative generics are not undermined by the existence of positive counterinstances. For example, a proponent of a generic like (81a) can coherently hold this sentence while recognizing that actually most Ks have an alternative property that is incompatible with F – that is, "crying". To explain this, Leslie argues that we need to postulate a polysemy for the referent of "Ks". For example, while generic (81a) refers to "ideal Ks" (normatively speaking), the following generic refers to "ordinary Ks" (descriptively speaking):

(82) Boys cry.

Normative generics are not limited to gender-specific sentences; for example, the following is a non-gender-specific normative generic:

(83) Friends don't lie.

What sentences like (81a), (81b), and (83) have in common is that they refer to "ideal Ks"; thus, they can be retained true despite several "ordinary Ks" are positive counterinstances. Although Leslie does not offer explicit truth-conditions for normative generics, she implicitly assumes that normative generics can be held true just in case a) they do not refer to "ordinary Ks" but to "ideal Ks" and b) F is a property associated with the "ideal Ks" by cultural norms. Thus, the truth of these sentences varies depending on the culture and the context we are considering.

Taken together, the different generics' truth-conditions proposed by Leslie can be summarized as follows (adapted from Leslie, 2008: 43):[15]

1. When "Ks" refers to "ordinary Ks", "Ks are F" is true just in case the counterinstances are negative, and:
(a) If F lies along a characteristic dimension for the Ks, then some Ks are F, unless K is an artifact or social kind, in which case F is the function or purpose of the kind K;
(b) If F is striking, then some Ks are F and the others are perceived to be disposed to have F;
(c) Otherwise, most Ks are F.

2. When "Ks" refer to "ideal Ks", "Ks are F" is true just in case F is a norm culturally associated with "ideal Ks".

[15] Leslie points out that these conditions are not part of the semantic content of generics; rather, they are the specifications of the circumstances under which we understand the proposition expressed by generics as true (see Leslie, 2007: 386 – 388; Leslie, 2008: 43 – 44).

Leslie's content-based proposal can account for our intuition on the truth and falsity of the troublesome generic sentences discussed so far. For example, consider again some of the sentences that have been used to discuss the flexibility of generics' truth-conditions in § 1.2.1.1:

(45b) Ravens are black.
(47b) Ducks lay eggs.
(49b) Ducks are female.
(50b) Mosquitoes don't carry the West Nile virus.
(51c) Humans are right-handed.
(51d) Humans are left-handed.

According to Leslie's theory, (45b) is correctly predicted to be a true characteristic generic inasmuch a) some ravens are black; b) "being black" lies along a characteristic dimension for animals (i.e., "color of the plumage") and, therefore, for ravens qua animals; and c) the non-conforming Ks – that is, the albino ravens – are negative counterinstances. Similarly, (47b) is correctly predicted to be a true characteristic generic; indeed, some ducks lay eggs, "laying eggs" lies along a characteristic dimension for ducks, and the non-conforming ducks simply lack the ascribed property. (49b), instead, is correctly predicted to be false; since being a male duck is incompatible with being a female duck, male ducks are a positive counterinstances to female ducks. Finally, Leslie's account handles our intuition about the falsity of generics like (50b), (51c), and (51d) as well. Despite most mosquitoes do not carry the West Nile virus and most humans are right-handed, (50b) and (51c) cannot tolerate the existence of positive counterinstances – the mosquitoes carrying the relevant virus and the left-handed humans, respectively. For the same reason, (51d) is undermined by the existence of positive counterinstances – that is, the right-handed humans.

Leslie's proposal can account for our intuitions on a great variety of generic sentences; moreover, she identified several cognitive factors involved in the interpretation of generics, which have been confirmed experimentally (see Chapter 2). Her account also provided

innovative tools to address the use of generics as stereotypes (see § 2.2.3). However, despite the explanatory power and the importance of her proposal, Leslie's account presents some theoretical issues.

A first problem is that there is no linguistic basis for the categorization between characteristic, striking property, majority, and normative generics. Since these types of generics have the same surface form, their categorization relies entirely on the subjective perception of the properties they ascribe. The consequence is that it is not possible to identify the truth-conditions of a given generic without knowing how F will be perceived. For example, while most of the times it is likely that sentences like (45b) or (47b) will be perceived as true characteristic generics, it is not clear how a generic like the following should be categorized and evaluated:

(84) Parasitoid wasps lay eggs in or on alive hosts.

(84) can be characterized both as a characteristic and a striking property generic. Indeed, the ascribed property lies along a characteristic dimension for parasitoid wasps (i.e., "mode of reproduction"); however, the ascribed property is also likely to be perceived as striking. It is possible to reply to this objection by observing that (84) can be considered true both as characteristic and as striking property generic; if this is the case, then (84) is only a non-problematic mixed case. On the contrary, I consider (84) problematic since it might be perceived false under one interpretation and true under another. For example, an entomologist could reject (84) as a striking property generic inasmuch only female wasps lay eggs. For the entomologist, wasps would not have a shared disposition to lay eggs in or on alive hosts – only female wasps would.[16] However, the entomologist can still accept (84) as a true characteristic generic. Indeed, according to Leslie, characteristic and striking property generics have different truth-conditions; in particular, the truth of characteristic generics do not require that the non-conforming Ks should be disposed to have F under the right

[16] For the same reason, an entomologist might reject "Mosquitoes carry the West Nile virus"; indeed, only female mosquitoes feed on blood and transmit diseases. Thus, male mosquitoes might not be a good predictor of "carrying the West Nile virus" for an entomologist.

circumstances. Furthermore, an entomologist – being used to the peculiar behaviors of certain insects – could not perceive (84) as striking at all. Thus, the truth-conditions of (84) might vary depending on our perception of F. This issue is not limited to (84). Consider the following sentences:

(85)

a. Mobsters shoot without hesitation.

b. Doctors save lives.

(85a) might be perceived both as a normative and as striking property generic; indeed, it could be used to express an in-group cultural norm, but also to describe the dangerous behavior of mobsters. Depending on the interpretation, (85a) will have different truth-conditions. Likewise, (85b) might be perceived both as a normative and as a characteristic generic. Under different interpretations, (85b) could express a norm concerning the "ideal doctors" or describe a function (i.e., a characteristic property) of doctors. In conclusion, Leslie offers a detailed and accurate report for the reasons why we judge a generic as true or false ex post facto. To establish the correct truth-conditions of a given generic, Leslie's account needs information on how the participants would perceive F, along with their background expectations, proximity with Ks, and expertise on Ks. Without this contextual information, the account cannot predict the truth-conditions of a given generic. As I discuss in Chapter 2, the lack of an obvious criterion for determining generics' truth-conditions makes it hard to contrast generic expressing stereotypes (Haslanger, 2011, 2014; Langton, Haslanger, & Anderson, 2012; see § 2.2.3).

Another problem concerning Leslie's characterization of negative and positive counterinstances has been recently raised by Nickel (2016). As he notes, Leslie's proposal correctly predicts both (45b) and (86) to be true:

(45b) Ravens are black.

(86) Cardinals are red.

Exactly as for (45b), according to Leslie, (86) is true inasmuch a) some cardinals are red; b) "being red" lies along a characteristic dimension for animals (i.e., "color of the plumage") and, therefore, for cardinals qua animals; and c) the non-conforming Ks – that is, the pale brown cardinals – are negative counterinstances. However, if this is the case, then (45b) entails that all the non-black plumage colors are negative counterinstances, whereas (86) entails that all non-red plumage colors are negative counterinstances. Consequently, if some ravens or cardinals have a color C, where C is any possible color, then the following sentences are true:

(87)

a. Ravens are C.

b. Cardinals are C.

However, it is intuitively false that the sentences in (87) are true for any possible C.[17] According to Nickel (2016), this issue can be avoided by conceiving the negativity of a property as kind-relative. This way, it would not follow from the fact that red is negative for ravens that the same color is also negative for cardinals. However, as I discuss in Chapter 2 (see § 2.1.2), Leslie's theory relies on the idea that the use and interpretation of generics reflect the functioning of a default mode of generalizing cognitive mechanism. According to Leslie, such a mechanism is fast and automatic. As Nickel notes, a mechanism characterized in this way could not discriminate between positive and negative properties differently for every kind.

In this chapter, I outlined the reasons why generics have been extremely debated in linguistics and philosophy of language. In the first part, I provided a classification of the linguistic phenomena that have been categorized as generics to clarify the focus of my research. In the second part, I described how generics differ from quantified and existential constructions. In this final part, I discussed two different approaches to the problem of

[17] As noted by Nickel (2016: 90 – 91), this problem is also highlighted by other pairs of sentences not color-related, like "Mice are short-lived" and "Turtles are long-lived".

generics' meaning. The fact that generics are outstandingly common in everyday speech and yet so complex to analyze within a theoretic framework makes them particularly interesting from the cognitive perspective. In recent years, an extensive literature regarding the relation between generics and cognition has developed. In the next chapter, I present some of the main experimental studies that have been conducted on the use and the interpretation of generics, with particular attention for the implications that these studies have for the analysis of social generics.

2. GENERICS, COGNITION, AND STEREOTYPING

In recent years, there has been a growing interest in experimental psychology for generics – particularly, for what their use and interpretation might disclose about the way organize our conceptual knowledge about categories. In this chapter, I discuss some of the main experimental works carried out on generics along with their cognitive implications. In the first part, I discuss the experimental data regarding the hypothesis that generics are default generalizations. In the second part, instead, I consider the influence of generics on social cognition, a topic that has been recently debated by both cognitive scientists and philosophers of language. Specifically, I focus on social generics' ability to promote essentialist beliefs and the implications of the content-based framework for stereotyping.

2.1 Generics as default generalizations

Much of what we learn depends on our ability to generalize. This ability allows us to extend the information we acquire beyond specific contexts and to make predictions every time we are dealing with known instances. As noted in the previous chapter, explicit generalizations can be expressed either through generics or quantified sentences, which present distinct linguistic properties. These linguistic differences also determine different cognitive implications for the two types of generalizations. In the present section, I discuss the hypothesis recently put forward by Gelman (2003, 2010) and Leslie (2007, 2008, 2012b) that generics are default generalizations, whereas quantified sentences lead us to generalize in non-default and specific ways. In the first part of this section, I provide an overview of the empirical data that motivates this view. In the second part, I present in more detail the hypothesis that generics are default generalizations. Finally, in the third part, I discuss the experimental evidence that indirectly confirms this hypothesis.

2.1.1 Motivations

The ability to generalize begins to manifest itself early in life. Infants under one year of age expect perceptually similar objects to share the same non-obvious properties (e.g., the production of a distinctive sound), but they rarely do when the objects are perceptually different (Baldwin, Markman, & Melartin;1993). Such expectations begin to be language-driven shortly after. For example, 13-month-old infants have the same expectations when two objects are introduced with the same labels, even when they are perceptually different (Graham, Kilbreath, & Welder, 2001). Later, through cognitive development and the acquisition of a complex language, we begin to comprehend and produce generics and quantified generalizations, which allows us to share knowledge without the need to repeat the relevant observations individually.

As noted in the previous chapter, generics are very hard to analyze. Despite all the work that has been done on their semantics, there is still no theory capable of fully accounting for their linguistic functioning. Conversely, quantified sentences have straightforward semantic rules based on the proportion of category members Ks displaying the property F. As Gelman (2003, 2010) and Leslie (2007, 2008, 2012b) noted, the theoretical intractability of generics becomes paradoxical when related to a wide range of empirical data. In what follows, I summarize the data they considered, together with some more recent findings.

First, the use of generics is markedly frequent in parent-child conversations (e.g., Gelman, 2003; Gelman, Coley, Rosengren, Hartman, Pappas, & Keil, 1998; Gelman, Taylor, & Nguyen, 2004). For example, Gelman et al. (1998) recorded the number of generics mothers produced while looking at picture books about categories with their 20- or 30-month-old children.[1] For both groups, in a session of 15 minutes, 3% to 4% of the mothers' utterances were generics on an approximate average of 189. «By extrapolation», as noted by Gelman (2003: 172), «this suggests that children would typically hear over 30 generics per hour, if

[1] It should be noted that this record included any possible generic nominal NP, not generic BP NPs only. For example, the utterance "A chipmunk's a little smaller than a squirrel" was identified by Gelman et al. (1998) as a generic. Although my research focuses on BP generics, in Chapter 1 I described how generics with nominal arguments can have a variety of syntactic forms, including IS (see § 1.1.).

placed in a comparable context». While maternal speech presented a high frequency of generics, quantified sentences were rare in the same interactions.

Second, children acquire generics at an early stage of development, starting around 30 months of age, and they use generics as often as adults by the age of 4 (e.g., Cimpian & Markman, 2008; Gelman, Goetz, Sarnecka, & Flukes, 2008; Gelman & Raman, 2003; Graham, Nayer, & Gelman, 2011). For example, Gelman and Raman (2003) found that children as young as 2- to 3-years-old begin to distinguish generics from non-generics in a way consistent with that of adults. In one of their studies, Gelman and Raman (2003) compared children's comprehension of generic questions, like (1a), to their comprehension of non-generic questions having a plural NP preceded by the determiner "the", like (1b):

(1)

a. Do Ks have F?

b. Do the Ks have F?

Participants were shown a series of realistic drawings, each depicting two Ks that were atypical or unusual in some respects. For example, one drawing depicted two penguins, which are atypical birds inasmuch they do not fly. Then, participants were asked a question related to the atypical property, either in generic or non-generic form. For example, while looking at the penguins drawing, participants were asked either (2a) or (2b):

(2)

a. Do birds fly?

b. Do the birds fly?

Gelman and Raman (2003) predicted that the absence or the presence of the determiner should signal whether the question is generic or non-generic, respectively. In turn, according to their prediction, participants should have answered in reference to the broad category in the first case and in reference to the depicted Ks in the second. In the case of the penguins

drawing, they should have answered "yes" to question (2a) and "no" to question (2b). Despite the salience of the atypical Ks, which created a context more inclined toward a non-generic interpretation, participants' responses in each age group (2-years-olds, 3-years-olds, 4-years-olds, and adults) were consistent with Gelman and Raman's (2003) prediction.

Third, Goldin-Meadow, Gelman, and Mylander (2005) found evidence that generic language is symbolically developed even in the absence of linguistic input. Examining videotapes of U.S. and Chinese deaf children that did not receive natural language input (due to both hearing losses and lack of exposure to formal sign languages), and that so spontaneously developed their own "home signs" system to communicate (Goldin-Meadow, 2003), Goldin-Meadow et al. (2005) observed gestures that can be traced back to generics. Furthermore, Goldin-Meadow et al. (2005) found that these children produce generics as frequently as hearing children in the respective countries. Goldin-Meadow et al.'s (2005) procedure was very conservative. To be classified as generics, the features described by the relevant gestures had to reflect general characteristics instead of idiosyncrasies. For example, if a deaf child reacted to a picture of an elephant wearing a bow by producing a "trunk" gesture, the same gesture could be initially classified as a generic. However, if the child produced a "bow" gesture instead, the gesture would not be classified a generic. Furthermore, the same features had not to be present in the context (e.g., as objects or in pictures); in the same example, the "trunk" gesture would not have been classified as a generic unless the relevant feature was not visible. Finally, the gesture had to refer to a class of objects instead of specific objects; for example, if the "trunk" gesture was accompanied by a negating gesture, then the gesture was interpreted as referring to the depicted elephant only.

Fourth, generics are acquired before quantified sentences; scholars found evidence for this testing native speakers of English (Hollander, Gelman, & Star, 2002), Mandarin Chinese (Tardif, Gelman, Fu, & Zhu, 2012), and Southern Peruvian Quechua (Mannheim, Gelman, Escalante, Huayhua, & Puma, 2010). For example, in a comprehension task, Hollander et al. (2002) found that while 4-years-olds and adults interpret generic questions differently from quantified questions, 3-years-olds treat both as generic questions. The following are examples of the questions used by Hollander et al. (2002):

(3)

a. Do animals eat?

b. Do all animals eat?

c. Do some animals eat?

Although 3-years-olds did not differentiate between generics and quantified sentences, their pattern of responses for generic questions was indistinguishable from the one displayed by older children and adults. To rule out the possibility that a lack of knowledge influenced their results, Hollander et al. (2002) checked whether children could distinguish between different types of properties: *wide-scope* (i.e., had by all or almost all members), *narrow-scope* (i.e., had by only some members), and *irrelevant* (i.e., had by none members). The sentences in (3) ascribe wide-scope properties. The sentences in (4) and (5), instead, ascribe narrow-scope properties and irrelevant properties, respectively:

(4)

a. Do bears have white fur?

b. Do all bears have white fur?

c. Do some bears have white fur?

(5)

a. Are children made of feathers?

b. Are all children made of feathers?

c. Are some children made of feathers?

Since children consistently distinguished among the three types of properties, Hollander et al. (2002) could rule out the potential influence of lack of knowledge.

Fifth, children interpret generics by adhering to some of the distinctive patterns presented by adults. In addition to evaluating generics as true despite salient exceptions (e.g., Brandone, Gelman, & Hedglen, 2015; Gelman & Raman, 2003), they recognize that generics are not

simply based on the proportion of Ks having F (e.g., Brandone, Cimpian, Leslie, & Gelman, 2012). For example, Brandone et al. (2012) found that similarly to adults, children accept generics ascribing characteristic properties of Ks holding for only one gender, like (6a), but reject generics ascribing non-characteristic properties of Ks holding for the same gender, like (6b):

(6)

a. Lions have manes.

b. Lions are boys.[2]

Conversely, children have difficulty mastering quantified sentences throughout childhood (e.g., Brandone et al., 2015; Brooks & Sekerina, 2006; Papafragou & Schwarz, 2006). For example, Papafragou and Schwartz (2006) observed that the misinterpretation of the quantifier "most" persists through preschool to middle childhood. In one of their studies, Papafragou and Schwartz (2006) used pictures where some fictional characters asserted that they were going to complete an action concerning other depicted objects. For example, in one picture there was a fictional character with six unlit candles, together with the following statement:

(7) I'm going to light most of these candles.

Then, participants looked at pictures where they had to evaluate whether the asserted action was completed. In the case of (7), they had to evaluate whether the assertion was true or false based on the number of lit candles (the possible ratios ranged from 0/6 to 6/6). Papafragou and Schwartz (2006) found that, contrary to adults, even 9- to 11-years-olds misinterpret the meaning of statements like (7) (i.e., they accepted ratios below 4/6).[3]

[2] Brandone et al. (2012) used the terms "boys/girls" instead of "male/females" due to the young age of the participants. A pilot test confirmed that children did non interpret these terms as referring to human boys and human girls.

[3] In a subsequent work, Halberda, Taing, and Lidz (2008) observed that preschoolers manifest a little competence with the quantifier 'most' if the task is made easier (e.g., telling whether most crayons in a box are yellow or blue). However, more recently, Sullivan, Bale, & Barner (2018) found evidence that such competence is far from the adult-like one. According to

Finally, while generics have been attested in every documented language (including pidgins; see Maurer & Meeuwis, 2013), no known language has a dedicated and explicit GEN marker (Dahl, 1975; Krifka, Pelletier, Carlson, ter Meulen, Chierchia, & Link, 1995). As noted in Chapter 1, while quantified sentences are marked by clear and explicit operators (i.e., the inclusion of quantifiers), generics' linguistic cues are varied and inconsistent (e.g., generics can have several syntactic forms, which can also be used to express existentials; see § 1.2). A fascinating case is that of the Amazonian language Pirahã, which would seem to lack quantifiers, but not generics (see Everett 2005, 2009).

The contrast between the theoretical analyses of generics and the empirical data presented above is clear. As noted by Gelman (2003, 2010) and Leslie (2007, 2008, 2012b), the linguistic complexity of generics should constitute a rather intricated learning puzzle for children. However, not only are generics markedly frequent in parent-child conversations; they are also early acquired, and their production is developed even in the absence of linguistic input. The paradox of the acquisition of generics is made even more evident by the fact that quantified sentences, which are decidedly less complex from a theoretical point of view, are acquired later than generics. To further complicate the issue, generics lack a straightforward set of markers, whereas quantified sentences present cohesive linguistic cues; consequently, generics should be more difficult to learn than quantified sentences.

To solve the contrast between generics' theoretical complexity and ease of acquisition, Gelman (2003, 2010) and Leslie (2007, 2008, 2012b) hypothesized that – contrary to quantified sentences – generics are default generalizations. In the next subsection, I consider their hypothesis in greater detail, together with the reasons why it would solve the learning puzzle of generics.

Sullivan et al. (2018), Halberda et al.'s observations (2008) would be due to a well-known early bias that leads children to select larger sets when interpreting unfamiliar quantifiers. In their view, children acquire actual competence with "most" around or later the age of 6.

2.1.2 The hypothesis

The wide empirical data considered in the previous subsection show that children acquire generics before quantified sentences and despite the absence of dedicated markers. To account for these observations, Gelman (2003, 2010) and Leslie (2007, 2008, 2012b) put forward a compelling hypothesis. In their view, children must have an early cognitive predisposition that allows them to elaborate generics. Then, increasing their linguistic competence and cognitive ability, children must learn to identify marked sentences as referring to specific proportions of individuals. In other words, Gelman (2003, 2010) and Leslie (2007, 2008, 2012b) hypothesized that generics are unmarked and easy for children because they are cognitively default generalizations; quantified generalizations, instead, are marked and hard for children because they are cognitively demanding, non-default, and sophisticated sentences. To illustrate this point, Leslie (2007, 2008) refers to Chomsky's (2000) discussion of automatic interpretations of other linguistic expressions, like the following:

(8) Jones climbed the mountain.

As Chomsky (2000) highlights, our default interpretation of (8) correspond to (9a) but not to (9b):

(9)

a. Jones climbed up the mountain.

b. Jones climbed down the mountain.

According to Chomsky (2000), while the interpretation (9a) is automatic, we need to use the explicit marker "down" to make the interpretation (9b) available. Similarly, according to Gelman (2003, 2010) and Leslie (2007, 2008), quantifiers serve as markers to generalize in sophisticated and specific ways, whereas generics – being unmarked – lead us to generalize

in a default manner. Specifically, Gelman (2003, 2010) hypothesized that the basic assumption for language users is that a statement about a category, like (10), is a generic:

(10) Dogs bark.

As Gelman notes (e.g., 2003), such assumption can be blocked by several semantic and pragmatic markers. For example, it can be blocked by the use of determiners, as in (11a); precise numbers, as in (11b); deictics, as in (11c); continuous and progressive aspects, as in (11d); and conventions of discourse, as the one allowing to resolve the anaphora of "they" in (11e):

(11)

a. The dogs bark.

b. Three dogs bark.

c. These dogs bark.

d. Dogs are barking.

e. These are the dogs of my neighborhood. They bark.

These markers serve to locate an utterance within specific boundaries and to identify which members of the category are concerned.[4] Contrary to specific utterances, generics express unbounded information (see Chapter 1, §1.1.4). According to Gelman (2003, 2010), children do not learn linguistic cues allowing them to *identify* generics; rather, they learn which elements might *block* a generic interpretation and suggest a specific one instead. Similarly, Leslie (2007, 2008, 2012b) hypothesized that the use and interpretation of generics reflect the functioning of a *default mode of generalizing* cognitive mechanism, whereas quantifiers direct our cognitive system to generalize in non-default and specific ways.

[4] Since my focus is on generics having BP NPs, I provided examples of markers that are relevant for this type of sentences. However, some of these markers are also relevant to generics having DS NPs, like "The dog barks" and IS NPs, like "A dog barks". Other cues, like tense and world knowledge, might be more relevant to these two types of generics. For example, the assumption that sentences having DS and IS NPs are generics can be blocked by the use of a past tense, as in "The dog barked", and the use of properties that are typically predicated of single individuals, as in "A dog is sick".

As noted by Leslie (2007), the hypothesis that generics – contrary to quantified sentences – are default generalizations is compatible with the broader dual-process theory of cognition (e.g., Kahneman & Frederick, 2002; Stanovich & West, 2000). According to this theory, two different cognitive systems – labeled System 1 and System 2, respectively – are involved in our reasoning processes. System 1 corresponds to the cognitive operations that are fast, effortless, intuitive, automatic, and whose processes are opaque – that is, whose processes function by means of pre-set cognitive rules we are not aware of. System 2, instead, corresponds to the cognitive operations that are slow, effortful, analytical, deliberate, and governed by transparent rules – that is, governed by processes we can accurately report. The two systems are suited to solve different orders of problems. For example, System 1 allows us to immediately recognize a specific emotion (e.g., fear) on a person's face and localize the source of a sound; System 2, instead, allows us to solve equations and determine the validity of logical reasoning. Sometimes, the two systems give conflicting answers when used to solve the same problem. Consider the following example (Kahneman & Frederick, 2002: 58):

> A bat and a ball cost $1.10 in total. The bat costs $1 more than the ball. How much does the ball cost?

We tend to solve problems like this by relying on the effortless automatic answer provided by System 1, which in this case is "10 cents". The process by which we obtained this answer is opaque; the default system somehow leads to separate the sum "$ 1.10" into "$ 1" and "10 cents" immediately. To get the right answer, instead, we need to rely on the effortful and slow System 2. Based on the explicit rules of algebraic reasoning, it is possible to verify that the ball costs 0.5 cents.

Leslie (2007) suggests that System 1 and System 2 are responsible for the elaboration of generics and quantified sentences, respectively. In her view, generics reflect the functioning of the fast, automatic, effortless, and cognitively basic System 1, which is not sensitive to considerations about quantities. For this reason, generics are both easy to master and easily

acquired. According to Leslie (2007), quantified sentences, instead, reflect the functioning of the slow, deliberate, effortful, and analytical System 2, which involves higher-level thought processes that are sensitive to considerations about quantities. For this reason, they are harder to master and acquired later. Interestingly, cases like the ball cost show that System 1 may interfere with problems that pertain to System 2 if the task is too demanding, or simply because System 1 is effortless. As mentioned in the previous subsection, Hollander et al. (2002) found that 3-years-olds do not differentiate between generics and quantified sentences, whereas they interpret generics in the same way older children and adults do. Compatibly with Leslie's suggestion (2007), they hypothesized that «[…] generics may be a kind of default for young children. When information-processing demands are high, children may ignore quantifiers and interpret a wide range of utterances as generic» (Hollander et al., 2002: 893).

The hypothesis that generics are default generalizations led to a number of predictions, which have been empirically investigated. In the next subsection, I present the empirical works that gave further support to this hypothesis.

2.1.3 Confirmations

The hypothesis that generics are default generalizations provides a compelling solution to the learning puzzle of these sentences; furthermore, it offers an explanation for both generics' ubiquity and unmarkedness. This hypothesis generated several empirical predictions, some of which have been experimentally confirmed.

The main support to the hypothesis that generics are default generalizations comes from a series of works showing that quantified sentences are often recalled as generics, but not vice versa (Gelman, Leslie, Gelman, & Leslie, 2019; Gelman, Leslie, Was, & Koch, 2015; Gelman, Tapia, & Leslie, 2016; Leslie & Gelman, 2012; Sutherland, Cimpian, Leslie, & Gelman, 2015). These works provided an important indirect confirmation of the hypothesis that generics are default generalizations. Indeed, if processing generics is easier than processing quantified sentences, then we should expect that quantified sentences are more

likely to be mistakenly processed as generics than the other way around. In turn, we should expect that people might form generic beliefs based on quantified sentences, and then recall these sentences as if they were generics. For example, in Leslie and Gelman's studies (2012), participants were shown a series of pictures about different animals, and were told a fact about each animal either in generic form, like in (12a), or in quantified form, like in (12b), (12c), and (12d):

(12)

a. Bears climb trees.

b. All bears climb trees.

c. Most bears climb trees.

d. Some bears climb trees.

After receiving a distractor task, participants were asked to recall the information provided while looking at the same pictures. Leslie and Gelman (2012) found that both adults and 4-years-old children tend to recall generics as generics, whereas quantified sentences – with the exception of some-quantified sentences – were frequently recalled as generics. The observation that some-quantified sentences were less likely to be recalled as generics is consistent with the fact that "some" is often interpreted as pertaining to a limited subset of Ks, due to a *scalar implicature* effect (e.g., Bott & Noveck, 2004; Levinson, 1983).[5] Furthermore, similarly to the developmental pattern previously observed by Hollander et al. (2002; see § 2.1.1), Leslie and Gelman (2012) found that 3-years-olds were more inclined to recall some-quantified sentences as generics than adults and 4-years-olds. Finally, both 3- and 4-years-olds were more inclined to recall all-quantified sentences as generics than adults.

[5] Scalar implicatures arise when a speaker uses a poorly informative expression like "some" despite the existence of more informative terms like "all". In these cases, the listener infers that the speaker did not use a more informative term because they were not in the position to use it: the speaker knows that not all Ks have F or does not know whether all Ks have F. As a consequence, despite the two quantifiers are not incompatible per se (e.g., if all Ks have F, then some Ks have F), a speaker utterance of sentences like "Some Ks have F" conveys that "Not all Ks have F" or that the speaker does not know whether all Ks have F (e.g., Bott & Noveck, 2004; Levinson, 1983).

Since generics and quantified sentences in English have different lengths, Leslie and Gelman (2012) had to rule out the possibility that the participants in their studies were simply shortening the sentences they were recalling. Specifically, they controlled whether participants showed a tendency to shorten the predicates, and whether participants recalled generics beginning with "no", like (13a), as positive generics, like (13b):

(13)

a. No spiders have nine legs.

b. Spiders have nine legs.

These controls confirmed that the observed tendency to recall quantified sentences as generics was not due to their different lengths, as participants showed no tendency to shorten the predicates or to recall generics beginning with "no" as positive generics. However, these were only indirect controls. Importantly, a direct confirmation that Leslie and Gelman's (2012) results were not due to the different lengths of generics and quantified sentences has been provided by Gelman et al. (2016). Indeed, Gelman et al. (2016) replicated Leslie and Gelman's (2012) results testing both adult and children native speakers of Spanish. In this language, generics, like (14a), and other quantified sentences, like (14b), have the same number of words:

(14)

a. Las jirafas tienen lengua morada ("Giraffes have purple tongues").

b. Muchas jirafas tienen lengua morada ("Many giraffes have purple tongues").

Since Gelman et al. (2016) observed the same effect with native speakers of Spanish, they provided a direct confirmation that participants were not simply shortening the sentences they were recalling.

Further support to the hypothesis that generics – contrary to quantified sentences – are default generalizations comes from the real-time study carried out by Meyer, Gelman, and

Stilwell (2011). Meyer et al. (2011) observed that, under time pressure, universally quantified sentences tend to be evaluated as if they were generics. If the hypothesis is correct, this finding should be expected; indeed, a correct evaluation of generics should be faster and cognitively less costly than a correct evaluation of quantified sentences, which would require to inhibit the default way of processing generalizations. Consequently, time pressure should decrease the accuracy in the evaluation of quantified sentences, leading the participants to adopt the readily and automatically available generic interpretation of the same sentences. In their studies, Meyer et al. (2011) asked their participants, all adults, to indicate by pressing two different keys whether several generics and universally quantified sentences on a computer screen were true or false. The NPs of both generics and quantified sentences were presented 2000 ms and in isolation before the predicates. This procedure permitted to eliminate a potential systematic increase in response time due to the NPs being longer for quantified sentences. Half of the participants were instructed to answer as fast as they could; the other half of the participants, instead, was told to take all the time they needed to provide their answers. Meyer et al.'s (2011) main interest was to compare responses to *wide-scope* sentences that should be accepted in generic form, like (15a), but rejected in quantified form, like (15b):

(15)

a. Dogs have four legs.

b. All dogs have four legs.

To avoid the possibility that participants simply learned that quantified sentences were always false to give their answers, Meyer et al. (2011) inserted *full-scope* sentences, which should be accepted both in generic form, like (16a), and quantified form, like (16b):

(16)

a. Giraffes have long necks.

b. All Giraffes have long necks.

The insertion of full-scope sentences also allowed to test a prediction – that is, that full-scope generics should be evaluated more quickly than full-scope quantified sentences, as the generic interpretation should be more accessible despite the equivalent acceptability of the two types of sentences. Finally, Meyer et al. (2011) also included *irrelevant-scope* sentences, which should be rejected both in generic form, like (17a), and quantified form, like (17b):

(17)

a. Squirrels have beaks.

b. All squirrels have beaks.

The insertion of irrelevant-scope sentences allowed to avoid the possibility that participants simply learned that generics were always true to give their answers; furthermore, this insertion allowed to exclude an alternative hypothesis – that is, the possibility that a potential time difference in evaluating universally-quantified wide-scope sentences, like (15b), and generic wide-scope sentences, like (15a), simply reflect a difference in processing the falsification of the former compared to the verification of the latter. If participants are slower in evaluating universally-quantified wide-scope sentences than irrelevant-scope sentences, then a potential time difference would not reflect the mere time difference required to falsify or verify a given sentence.

Meyer et al.'s (2011) results were consistent with the hypothesis that generics are default generalizations. First, they found that speeded participants were more likely to evaluate universally-quantified sentences as if they were generics. Second, when speeded participants correctly rejected universally-quantified wide-scope sentences, they needed more time to do that than verifying both generic wide-scope sentences and universally-quantified irrelevant-scope sentences. Finally, both speeded and un-speeded participants tended to accept both universally-quantified full-scope sentences and generic full-scope sentences; however, as predicted, participants needed more time to accept universally-quantified full-scope sentences than generic full-scope sentences (even though they had equivalent acceptability). Meyer et al. (2011) hypothesize that this discrepancy is due to the fact that participants tried

to consider plausible counterexamples to the universally-quantified full-scope sentences before accepting them, but relied on the default generic representation to accept their generic counterparts. These results are appealing, but further research is required for their confirmation. For example, as Meyer et al. (2011) noted, possible further studies should expand the comparison between generics and quantified sentences by varying both the quantifiers and the content of the predicate being analyzed (e.g., testing most-quantified sentences and dangerous properties). In a later study, Lazaridou-Chatzigoga & Stockall (2013) provided evidence that the finding that generics are easier to process replicates with generics concerning highly prevalent characteristic dimensions of Ks, like (15a), but not with generics concerning characteristic dimensions held only by a minority of Ks, like (18):

(18) Snakes lay eggs.

Finally, possible support to the hypothesis that generics are default generalizations comes from the observation of a tendency to accept false universal generalizations when they would be true in generic form in case the ascribed property is characteristic (Leslie, Khemlani, & Glucksberg, 2011; Khemlani, Leslie, Glucksberg, & Rubio Fernandez, 2007). Leslie et al. (2011) called this overgeneralization from the acceptance of a characteristic generic to the acceptance of a corresponding quantified sentence the *generic overgeneralization* effect, or GOG effect. The GOG effect is compatible with the hypothesis that generics are default generalizations; indeed, if evaluating quantified sentences require to override our intuitive judgments, then we should expect that, in certain circumstances, we fail to override those judgments and interpret quantified sentences as if they were generics. In Leslie et al.'s (2011) studies, participants, all adults, were asked to accept or reject sentences that were presented either in generic form, like (19a), or quantified form, like (19b) and (19c):

(19)

a. Tigers have stripes.

b. All tigers have stripes.

c. Some tigers have stripes.

These sentences involved several types of predications: *quasi-definitional* (i.e., true for all Ks), as in (20a), *majority characteristic* (i.e., highly prevalent characteristic dimension of Ks), as in the sentences in (19) or in (20b), *minority characteristic* (i.e., only held by a subset of Ks but characteristic), as in (18) or in (20c), *majority* (i.e., highly prevalent but non-characteristic), as in (20d), *striking* (i.e., held by few Ks but perceived as dangerous), as in (20e), and *false generalization* (i.e., highly prevalent but rendered false by the presence of a salient alternative property among Ks), as in (20f):

(20)

a. Triangles have three sides.

b. Ravens are black.

c. Ducks lay eggs.

d. Cars have radios.

e. Pit bulls maul children.

f. Canadians are right-handed.

Quasi-definitional and false generalization predications were included for comparison reasons; among Leslie et al.'s (2011) items, the former were the only true sentences in universal form, whereas the latter were the only false sentences in generic form.

Leslie et al. (2011) observed a GOG effect when the predications were characteristic. In other words, they found that false universal generalizations ascribing characteristic properties, like (19b), are accepted in case the corresponding generics, like (19a), are true. Specifically, although they were all false generalizations, universally-quantified majority characteristic predications had an acceptance ratio of 78% and universally-quantified minority characteristic predications had an acceptance ratio of 51%. The GOG effect was not observed for the other types of predications. Leslie et al. (2011) hypothesize that such discrepancy is due to the fact that people are used to reasoning about characteristic properties as normally held by all the members of a kind. In these cases, there is a low risk of making an error of judgment; thus, participants might have avoided cognitive efforts by relying on

the default generic interpretation. Non-characteristic properties, instead, are not likely to be normally held by all the members of a kind. In these cases, there is a high risk of making an error of judgment; thus, participants might be aware that they should not rely on their intuitions when evaluating non-characteristic properties.

Leslie et al. (2011) conducted other studies to exclude three alternative explanations for this finding. The first alternative explanation is that the GOG effect is actually elicited by quantifier domain restriction (see Chapter 1, § 1.2.1.2). Accordingly, participants might have interpreted universally-quantified sentences as generalizing over sub-sets of Ks. For example, a sentence like (21a) might have been accepted because interpreted as (21b):

(21)

a. All ducks lay eggs.

b. All fertile female ducks lay eggs.

To rule out this alternative, Leslie et al. (2011) conducted a study where participants were primed to consider all existing individual Ks to avoid a restricted interpretation of the quantifier. For example, instead of (21a), participants were asked the following:

(22) Suppose the following is true: there are 431 million ducks in the world. Do you agree with the following: all ducks lay eggs.

The GOG effect was observed even when participants evaluated sentences like (22), although the effect was reduced. The second alternative explanation is that participants interpreted universally-quantified sentences as generalizing over sub-kinds. According to this alternative, a sentence like (21a) was accepted because interpreted as (23):

(23) All kinds of ducks lay eggs.

To rule out this second alternative, in a subsequent study, participants were asked to provide paraphrases of sentences like (21a) to check whether they used a language suggesting

a sub-kind interpretation. Sentences referring to sub-kinds were rarely present in the paraphrases; conversely, Leslie et al. (2011) observed a tendency to paraphrase universally-quantified sentences as generics.

The third alternative explanation to the GOG effect is that participants simply ignored that the universally-quantified characteristic sentences were false; for example, they ignored that only female ducks lay eggs and thus judged sentences like (21a) as if they were true. To rule out this possibility, participants were asked to complete the same task of the main study, but also to accept or reject false characteristic counterparts to the main items, like the following:

(24) Male ducks lay eggs.

Participants correctly rejected sentences like (24); however, the GOG effect persisted despite the acknowledgment of the existing counterexamples. Thus, Leslie et al. (2011) concluded that this alternative explanation could be ruled out as well.

The results provided by Leslie et al. (2011) are appealing. However, Lazaridou-Chatzigoga, Katsos, and Stockall (2017) and Lazaridou-Chatzigoga, Stockall, and Katsos (2019) recently provided evidence that quantifier domain restriction might be a viable alternative explanation for Leslie et al.'s (2011) results. In their view, the studies carried out by Leslie et al. (2011) leave open the possibility that the GOG effect is actually elicited by quantifier domain restriction. In particular, priming the participants to consider all existing Ks using items like (22) does not make the relevant exceptions salient (e.g., male ducks). Consequently, Lazaridou-Chatzigoga et al. (2017) and Lazaridou-Chatzigoga et al. (2019) investigated whether making the exceptions salient would limit the tendency to accept false characteristic universal generalizations, like (19b), when they would be true as generics, like (19a):

(19b) All tigers have stripes.
(19a) Tigers have stripes.

Lazaridou-Chatzigoga et al. (2017) and Lazaridou-Chatzigoga et al. (2019) presented participants with majority characteristic sentences after describing three different contexts: neutral, contradictory, or supportive. The following are examples of the three types of contexts:

Neutral context: Linton Zoo is home to three tigers, Tibor, Baginda and Kaytlin, whose playful games visitors love to watch and photograph.

Contradictory context: Linton Zoo is home to three tigers, Tibor, Baginda and Kaytlin, whose fur is all white due to a recessive gene that controls coat colour.

Supportive context: Linton Zoo is home to three tigers, Tibor, Baginda and Kaytlin, whose black and orange coats visitors love to photograph.

Lazaridou-Chatzigoga et al. (2017) and Lazaridou-Chatzigoga et al. (2019) found that the contradictory contexts strongly limited the acceptability of false characteristic universal generalizations. For this reason, they conclude that quantifier domain restriction might account for a significant proportion of what Leslie et al. (2011) consider GOG effect. However, these findings are also compatible with an alternative explanation that would reinforce the GOG effect hypothesis instead (see Karczewski, Wajda, & Poniat, 2020). Contradictory contexts like the one used by Lazaridou-Chatzigoga et al. (2017) and Lazaridou-Chatzigoga et al. (2019) might elicit a more careful analysis of the relevant generalizations compared to neutral or supportive contexts. In contradictory contexts, we might deviate from the GOG bias because we are more prompted to rely on analytical reasoning. Conversely, in neutral or supportive contexts, we might be more likely to accept false characteristic universal generalizations because these contexts do not inhibit our default way to overgeneralize the relevant sentences.

In conclusion, more research seems to be required on the specific GOG effect to assess the actual strength of its support to the hypothesis that generics are default generalizations.

Nevertheless, the overall empirical confirmations of this hypothesis, together with the data that motivates it, remain overwhelming. The works discussed in this section highlight the basic role that generics play in our cognition. The relevant experimental literature is progressively showing how generics are an early, immediate, and default way for transmitting beliefs about categories. For this reason, generics might deeply affect how we perceive such categories, including social groups. In the next section, I consider some of the theoretical and experimental works regarding the influence of social generics on our cognition.

2.2 Generics and social cognition

The relation between generics and social cognition has attracted both cognitive scientists and philosophers of language. Indeed, such relation has important implications for phenomena such as stereotyping and discrimination. The use and interpretation of generics seem to be highly related to *essentialism*, a cognitive bias that leads to believe that the members of certain categories, or Ks, share properties and dispositions due to an underlying nature, or essence (e.g., Gelman 2003; Medin and Ortony 1989). Specifically, empirical evidence shows that generics play a powerful role in leading both adults and children to develop essentialist beliefs about categories (e.g., Gelman, Ware, & Kleinberg, 2010; Rhodes, Leslie, & Tworek, 2012). At the same time, generics are a common means to describe essentialized categories (e.g., Gelman et al., 2004; Rhodes et al., 2012; I discuss this twofold relation in § 2.2.2). Importantly, these findings can be linked to Leslie's content-based approach to generics' meaning (e.g., 2007, 2008; see Chapter 1, § 1.3.3) to identify other specific ways in which these sentences affect our cognition. In her view, generics ascribing a property F perceived as having high informational value are easily accepted, regardless of the actual proportions of Ks displaying F. For this reason, such generics could distort the perception of Ks in a way that is difficult to challenge, with particularly harmful effects on the perception of certain social categories (e.g., Leslie, 2017).

In the present section, I primarily focus on the role that generics play in shaping our perception of social categories. Correspondingly, I mainly consider their relation to *social essentialism*, which specifically concerns essentialist beliefs about groups of people (Rhodes et al., 2012). In § 2.2.1, I provide a brief overview of what social essentialism is, and how it can affect the perception of social categories. In §2.2.2, I present the empirical evidence concerning the interplay between social generics and essentialist beliefs. In §2.2.3, I discuss the proposals concerning how the combination of Leslie's concept-based account with these findings highlights additional specific ways in which social generics might be harmful. Finally, in § 2.2.4 I consider some controversies linked to these proposals and findings.

Before proceeding, I need to point out that while generics might strongly facilitate the promotion of essentialist beliefs, other linguistic construction can also lead to essentialize the Ks they are about. For example, essentialist views can be also transmitted by the use of category labels (e.g., Brandone & Gelman, 2009; Waxman, 2010), quantified sentences (e.g., Cimpian & Erickson, 2012; see §2.2.4), and explanations appealing to category membership, like "she does that because she is a girl" (e.g., Gelman, Cimpian, & Roberts, 2018; Vasilyeva & Lombrozo, 2020). However, in this section I specifically focus on social generics and, to a lesser extent, on quantified sentences.

2.2.1 The harmful consequences of social essentialism

In the psychological literature, the term "essentialism" indicates a cognitive bias that leads to believe that certain Ks share properties and dispositions due to an underlying essence (e.g., Gelman 2003; Medin and Ortony 1989). Psychological essentialism should not be confused with the traditional notion of "essentialism" in the philosophical literature. The traditional notion refers to the set of metaphysical theories according to which the properties of an object are divided into essential (or necessary) and accidental (or contingent) properties (e.g., Cartwright, 1968). In psychology, the term "essence" has no metaphysical import; rather, it denotes the presumed existence of something hidden, internal, inherent, stable, and innate that is causally responsible for the properties of Ks. In turn, this essence would make such

Ks fundamentally similar to each other and, at the same time, fundamentally different from the members of other categories (e.g., Gelman & Markman, 1986; Medin & Ortony, 1989). Importantly, the psychological notion of "essence" has been described as a *placeholder* (e.g., Gelman, 2003; Medin & Ortony, 1989): one might believe that certain Ks have an essence without entertaining any explicit representation of what such essence is. Alternatively, one might form a representation of Ks' essence relying on the cultural notions of other hidden substances having a presumed causal force; such substances might be material, like Ks' blood or DNA, or non-material, like Ks' vital force and soul (Gelman, 2003). The essence placeholder grants broad inductive inferences about essentialized Ks, whether one has a clear representation of it (e.g., DNA) or not (e.g., something mysterious).

Essentialist reasoning begins to manifest itself around the age of 3 (Gelman, 2003). Evidence suggests that this bias emerges from an *inherence heuristic* – that is, an early intuitive tendency to prefer inherent explanations for the properties of certain Ks that, in turn, would lead to positing an internal essence responsible for these properties (Cimpian & Salomon, 2014; Salomon & Cimpian, 2014). The influence and pervasiveness of essentialism in children's expectations about categories have been extensively investigated and attested (e.g., Brandone, 2017; Gelman & Markman, 1986, 1987; Gelman and Wellman, 1991; Rhodes & Gelman, 2009).

While essentialist reasoning about animal categories might facilitate learning processes (e.g., Gelman & Markman, 1987), social essentialism – which is specifically used to describe essentialist reasoning about social categories (Rhodes et al., 2012) – often has harmful consequences. Indeed, social essentialism entails the belief that certain ways to categorize people, for example based on race, gender, sexual orientation, ethnicity, and so on, reflect objective, meaningful, and natural distinctions. By leading to believe that distinct social categories mark distinct kinds of people (e.g., Rhodes et al., 2012; Rhodes & Mandalaywala, 2017), this bias provides a solid cognitive basis for stereotyped beliefs and discrimination (e.g., Allport, 1954). Essentialized social categories tend to be the targets of prejudice and negative social attitudes (e.g., Haslam, Rothschild, & Ernst, 2000; Rhodes et al., 2012). Social essentialism also leads to support policies enhancing the boundaries between groups

(Roberts, Ho, Rhodes, & Gelman, 2017) and to promote the endorsement of social hierarchies as well as prejudice toward "lower-status" groups (Mandalaywala, Amodio, & Rhodes, 2018).

Importantly, essentialist beliefs affect our perception of and behavior toward the relevant Ks, whether we belong to these groups or not. For illustration purposes, let us briefly consider some of the consequences of essentialist beliefs about gender (of course, this illustration is not intended to be exhaustive). An essentialist view of gender entails the perception of men as being fundamentally all alike while being fundamentally different from women, and vice versa. People holding essentialist beliefs about gender tend to think on this basis that men have different personal traits, interests, and responsibilities than women (e.g., Coleman & Hong, 2008; Frable & Bem, 1985; Gaunt, 2006; Martin & Parker, 1995). Essentialist views about gender begin to form between 3 to 4 years of age (Rhodes, Gelman, & Karuza, 2014). Among the most pernicious gendered stereotypes, there is the widespread belief that women and men have different intellectual abilities (e.g., Bennett, 1996). This and other gender-related beliefs have detrimental consequences. For instance, by age 6, girls are less likely than boys to believe that the members of their own gender are "really, really smart" and, correspondingly, they start to avoid activities believed to be for children who are "really, really smart" (Bian, Leslie, & Cimpian, 2017). Later in life, women find themselves underrepresented in academic fields whose practitioners believe that "natural brilliance" belongs to men but not women (Leslie, Cimpian, Meyer, Freeland, 2015). Correspondingly, describing professional and educational opportunities as requiring brilliance for success undermine women's interest in them (Bian, Leslie, Murphy, Cimpian, 2018). Furthermore, in STEM majors, women have a higher abandon rate than men after failing an introductory calculus course, and a possible explanation for this is that gendered stereotypes about STEM fields cause different self-assessments in men and women after a poor performance (Sanabria & Penner, 2017).

Reasoning about social categories in essentialist terms is not inevitable. For example, Rhodes and Gelman (2009) found that the way we think about social categories is affected both by development and cultural context. Specifically, they found that 5- to 7-years-olds

view gender categories as natural kinds, but race categories as more conventional. However, as the age of the participants in their studies increased, they observed that children living in a midsize, liberal, and ethnically diverse community begin to represent both gender and race as conventional kinds; children growing up in a rural, conservative, and ethnically homogeneous community, instead, begin to represent both gender and race as natural kinds. Thus, contextual and cultural inputs, including the way we speak about certain categories, might be important factors in the development of social essentialism.

The transmission of social essentialism has been strongly associated with generic language. In the next part, I present the relation between generics and social essentialism, which has been widely investigated in the psychological literature.

2.2.2 Generics and social essentialism: a "self-reinforcing cycle"

The particular relation between generics and essentialism has been recently described as a "self-reinforcing cycle" (Vasilyeva & Lombrozo, 2020). On the one hand, essentialized categories are commonly described through generic language (e.g., Gelman et al., 2004; Rhodes et al., 2012; Segall, Birnbaum, Deeb, & Diesendruck, 2015). On the other hand, generics lead to essentialize the categories they are about (e.g., Gelman et al., 2010; Rhodes et al., 2012). This loop-like relation between generics and essentialism is supported by robust evidence. In what follows, I mainly focus on the studies carried out by Rhodes et al. (2012). There are two reasons for this choice. First, they tested social generics; consequently, their results are more relevant to the study of the transmission of social essentialism. Second, they tested both aspects of the "self-reinforcing cycle" – that is, 1) they examined the power of social generics to transmit essentialist beliefs, on the one hand, and 2) whether essentialized social categories are more frequently described using generics than non-essentialized categories, on the other hand.

To investigate the first aspect, Rhodes et al. (2012) conducted two studies. Their methodology is based on previous experiments carried out by Gelman et al. (2010) on the interpretation of generics about animals. To conduct their first study, Rhodes et al. (2012)

created three illustrated storybooks regarding a fictional group of people called "Zarpies". Every picture was accompanied by a text-line, which described the depicted property, involving generics, as (25a), or a specific language, with or without labels, as, respectively, (25b) and (25c):

(25)

a. Look at this Zarpie! Zarpies eat flowers.

b. Look at this Zarpie! This Zarpie eats flowers.

c. Look at this one! This one eats flowers.

To avoid the influence of possible participants' pre-existing essentialist biases, Zarpies were depicted heterogeneous for ethnicity, gender, and age. The participants of the first study were both adults and four-years-old children. They were randomly assigned to either the generic or one of the specific conditions.

Adults read the storybook twice; afterward, they were given a battery of tests designed to assess the extent to which they essentialized Zarpies. The acquisition of essential beliefs was measured by assessing the extent to which the participants expected a particular property to be innate, extensible to other members, or due to the category membership. Participants that took the possession of a property to be due to intrinsic causes, to pass on to offspring, and to be shared by other Zarpies, were interpreted as bearing essentialist beliefs. Participants that, instead, took the possession of a property to be due to incidental causes, to derive from upbringing, and not to be extensible to other Zarpies, were interpreted as not bearing essentialist beliefs.[6] Children, instead, listened to the storybook (read by an experimenter) two times in a first research session, then two times again in a second research session approximately three days later. After about three other days, they completed the tests measuring their essentialist beliefs about Zarpies. The first study confirmed that, as

[6] For example, participants were asked to predict whether a baby Zarpie raised by a non-Zarpie mother grew up having a) a property of the Zarpie mother or b) a property of the non-Zarpie mother. The first prediction was coded as an essentialist answer, the second one as a non-essentialist answer.

hypothesized by Rhodes et al. (2012), both adults and children in the generic condition essentialize Zarpies to a greater extent than the participants in the other two conditions.

The reason why participants in the generic condition essentialized Zarpies to a greater extent may not have been due to the use of generics per se; rather, the results of the first study may have been caused by the comparison between plural and singular NPs (see the previous chapter). Indeed, while the NPs occurring in the generic condition were plural, those in the two specific conditions were singular. To exclude this alternative hypothesis, Rhodes et al. (2012) conducted a second study employing generics with IS NPs as well. The storybook in the generic condition contained sentences with IS NPs, like (26a); the one in the specific condition was written with sentences like (26b):

(26)

a. Look at this Zarpie! A Zarpie sleeps in tall trees.

b. Look at this Zarpie! This Zarpie sleeps in tall trees.

This time, the study was completed in one research session for both adults and children. The book was read only twice to the children, and they completed the essentialist beliefs tests on the same day. To rule out that eventual differences were due to this change in the procedure rather than to the singular NP, a group of children was assigned again to a BP generic condition. Both adults and children in the generic conditions essentialized Zarpies to a greater extent than the participants in the specific condition. Furthermore, no difference was found between responses in the BP and IS generic conditions.

This study provided evidence that the results of the first experiment were not simply due to the comparison between plural and singular NPs, but that the generic construction by itself led to essentialize Zarpies more than specific statements. It also showed that the use of generics facilitated children's acquisition of essentialist beliefs very rapidly, while the first study demonstrated that this effect persists over time.

As for the second aspect, Rhodes et al. (2012) conducted a third and final study to test the hypothesis that, in parent-child interactions, holding essentialist beliefs about a social

category leads parents to produce more generics in describing that category. Parents were randomly divided into two groups. Zarpies were presented as a very distinct social category, both biologically and culturally, to the first but not the second group of parents. The text presented to the latter, instead, described Zarpies as a non-distinct social group. Subsequently, both groups received a book about Zarpies that contained the same pictures used in the previous studies, but without text-lines. Parents were asked to describe to their child the people and events depicted. The experimenters found that parents in the first group, who read the paragraph describing Zarpies as a very distinct social group, were more likely to produce generic language than those in the second group, who read the other paragraph. Finally, consistently with other studies in the literature showing that essentialized categories tend to be the targets of prejudice and negative social attitudes (e.g., Haslam et al., 2000), parents in the first group also produced more negative evaluations about Zarpies.

Rhodes et al.'s (2012) findings show that generics play a powerful role in leading both adults and children to develop essentialist beliefs about social categories. In turn, generics might profoundly foster the development of social discrimination and prejudice. For example, consider the following social generic:

(27) Girls are bad at math.

Given that generics facilitate the development of essentialist beliefs, accepting (27) might lead children to believe that girls are importantly alike and, at the same time, different from boys. If so, they would feel justified to judge a newly encountered girl by extending their knowledge about other girls. Moreover, entertaining essentialist beliefs about girls, they would believe that many features of individual girls depend on their membership to the relevant category. For example, they would believe that a girl's inability at math does not depend on her personal attitudes. Rather, they would believe that this inability derives from the underlying nature of the category, with all the harmful consequences that this entails (see the previous subsection). Correspondingly, given that generics and social essentialism seem

to promote each other in a self-reinforcing cycle, the acquisition of essentialist beliefs about gender would encourage the use of generics to describe the relevant category and vice versa.

Following Rhodes et al.'s (2012) studies, other works provided further evidence for these findings and expanded our understanding of generics' power to shape our social cognition (e.g., Foster-Hanson, Leslie, & Rhodes, 2016, 2019; Leshin, Leslie, & Rhodes, 2020; Rhodes, Leslie, Bianchi, & Chalik, 2018a; Rhodes, Leslie, Saunders, Dunham, & Cimpian, 2018b). For example, using a similar methodology to Rhodes et al. (2012), Rhodes et al. (2018a) found that generics play a critical role in the development of social categorization compared to specific sentences, but also found that this extends to the comparison between generics and specific sentences having a plural form, like (28):

(28) These Zarpies eat flowers.

Rhodes et al. (2018a) did not measure *essentialism* as the dependent variable. Nonetheless, together with Rhodes et al.'s (2012) findings, their results further support the hypothesis the impact of generics on our social cognition is linked to the use of the generic construction per se, rather than being simply due to the comparison between plural and singular NPs. Foster-Hanson et al. (2016, 2019), instead, found that generics about Zarpies, like (29a), facilitate the transmission of social essentialism even when they are later negated, as in (29b), replaced, as in (29c), or rejected, as in (29d):

(29)

a. Zarpies have striped hair.

b. No, no, no! Zarpies don't have striped hair.

c. No, no, no! Zarpies have spotted hair.

d. No, no, no! That's not right about Zarpies.

Specifically, Forster-Hanson et al. (2016, 2017) found that the generic corrections in (29) lead children to essentialize Zarpies to a greater extent than specific corrections, like (30):

(30) No, no, no! This Zarpie has striped hair.

These results illustrate an additional subtle danger of social generics: using generic language to challenge others' use of social generics might still transmit the belief that the relevant categories have a common essence. As these works indicate, negating or rejecting social generics like (27) using sentences like (31a), (31b), or (31c) might not constitute an effective strategy to reduce the transmission of essentialist beliefs about girls:

(31)

a. Girls aren't bad at math.

b. Girls are good at math.

c. That's not right about girls.

The sentences in (31) might still convey that girls have a common essence. Even generics ascribing positive properties like (31b), which might not seem intuitively problematic, could promote essentialist beliefs about girls, just like (27) does. On this basis, children might be more likely to accept generalizations over girls. Consequently, at a later time, they might paradoxically come to accept sentences like (27) based on the belief that girls have a presumed common essence.

This possibility is not far-fetched, as illustrated by several works. For example, Kay, Day, Zanna, and Nussbaum (2013) found evidence indicating that this is precisely what happens: a positive stereotype makes it easier to accept a negative one. Specifically, Kay et al. (2013) found evidence that positive stereotypes are more likely to induce essentialist beliefs than negative ones and that, in a later time, they make people more willing to attribute negative properties to the relevant categories. Kay et al. (2013) hypothesized that positive stereotypes evade red flags that, instead, block the acceptance of negative stereotypes; consequently, the former are more easily accepted than the latter. For this reason, Kay et al. (2013) conclude that positive stereotypes can even be more harmful than negative ones.

Another example comes from Chestnut and Markman (2016), which found evidence that directional statements like (32),

(32) Girls can do math as well as boys.

which are commonly used to express gender equality, may conversely reinforce the idea that boys set the standard, and that being good at math is atypical for girls.

Finally, Cimpian (2010) and Cimpian, Mu, and Erickson (2012) found that gender-related generics can be detrimental for children even when they ascribe positive features. In their studies, children were asked to participate in an unfamiliar game. The game description contained a gender-related generic, like (33a) or a specific sentence, like (33b):

(33)

a. Boys [girls] are really good at this game.

b. There is a boy [girl] who is really good at this game.

Children in the generic condition displayed an impaired motivation (Cimpian, 2010) and performance (Cimpian et al., 2012) than children in the specific condition, regardless of whether they belonged themselves to the relevant category. As a possible explanation, Cimpian et al. (2012) hypothesized that generics like (33a) lead children to think that they have little control over their performance. This hypothesis is based on previous studies carried out by Cimpian and Markman (2011), which shows that children tend to explain the facts described through social generics like (33a) as due to stable features that effortlessly emerge from an internal source of the category members. Interestingly, and compatibly with the other works considered here, Cimpian and Markman (2011) also found that when the same abilities are described using specific sentences, like (33b), or generics about more context-based categories (e.g., "Boys [girls] at this school"), children tend to explain the relevant facts in non-essentialized ways. For example, they trace back Ks' abilities to practice rather than something linked to Ks' internal source.

In conclusion, extensive evidence suggests that social generics are problematic. By leading to essentialize the categories they are about, they might expose the relevant categories to future discrimination and prejudice, even if they ascribe apparently good features. However, the evidence discussed in this subsection also suggests that using a more specific language or narrowing the focus on more context-based categories might prevent the transmission of essentialist beliefs and be used to correct generic statements. Nevertheless, generics ascribing properties that are particularly relevant for our conceptual knowledge might further complicate the issue. According to Leslie's content-based approach to generics' meaning (e.g., 2007, 2008), one might accept a generic ascribing a property perceived as having high informational value despite being aware that almost all Ks do not have it. For example, social generics ascribing properties like "being terrorists" might be more readily accepted but also harder to challenge than generics ascribing less salient properties, worsening their impact on social cognition. In the next part, I present the proposals based on Leslie's theory highlighting how social generics might make prejudicial beliefs even harder to refute.

2.2.3 The content-based approach and the pitfalls of social generics

Generics are accepted even based on weak evidence. For this reason, they can be held even in face of many exceptions (see Chapter 1, § 1.2.1.1). This feature can make generics especially harmful when used to describe social categories. Indeed, a person believing a generic generalization will be unlikely to change their belief if contradicted. As Gelman observes (2003: 174), the resistance to exceptions displayed by generics makes them particularly difficult to challenge: «Even erroneous properties stated generically, such as stereotypes concerning gender or race, may be more difficult to counter and erase than erroneous properties stated absolutely».

According to Leslie's content-based approach (e.g., 2007, 2008; see Chapter 1, § 1.3.3), not all generics tolerate exceptions in the same way. In her view, the interpretation of generics is determined by the informational value of the properties they predicate. More precisely,

Leslie claims that the interpretation of generics reflects how the default cognitive mechanism presented in this chapter deal with the informational value of the ascribed properties. In particular, according to Leslie (2007, 2008, 2012b), striking property generics reflect how the default cognitive mechanism deals with particularly dangerous properties, granting generalizations based on spotty evidence – provided that Ks are good predictors of the property in question (see the previous chapter).[7] Characteristic generics, instead, are a linguistic outlet of the way this mechanism exploits regularities in the world, leading us to characterize kinds according to different dimensions (e.g., characteristic diets, noises, modes of locomotion, and modes of reproduction). Finally, majority generics reflect how this mechanism deal with information that is not distinctive or dangerous: if most Ks have the relevant neutral properties, we accept the correspondent generic or reject it otherwise.

Since the default cognitive mechanism deals with striking and characteristic generic predications in a matter that is not based on the actual proportion of Ks displaying F, Leslie (2007, 2008) hypothesizes that the relevant generics are more easily accepted than generic predications without such connotations. In particular, properties that are perceived as distinctive or dangerous have relatively high informational value, and, in turn, might be more prominent in our conceptual knowledge, and so more readily accepted in generic form. This hypothesis, if correct, has important consequences for the interpretation of social generics. For example, consider the following sentence:

(34) Muslims are terrorists.

Several scholars (e.g., Langton, Haslanger, & Anderson, 2012; Leslie, 2017) hold that social generics ascribing dangerous properties like (34) are especially pernicious because they could be accepted more easily than social generics with different valences. For example, Leslie (2017: 399) argues that «just as it takes but a few instances of sharks attacking bathers

[7] Interestingly, as Leslie (2007) notes, her view is compatible with other works showing that System 1's evaluations are sensitive to the strikingness of the information provided and depend on how the relevant information is presented. For example, testing choices based on mortality and survival rates, McNeill, Pauker, Sox, and Tversky (1982) found that people tend to prefer a medical procedure having a 90% short-term survival rate than a procedure having a 10% immediate mortality rate, as the latter is perceived as more threatening.

[…] for us to make the corresponding category-wide generalization, so also a strikingly negative action of a few members of a racial, ethnic, or religious minority may lead others to form a general belief concerning their entire group». According to Leslie, our default cognitive mode of generalizing mechanism, manifested in language through generics, is biased to extend threatening behaviors of few individuals to salient categories as a whole, regardless of the actual prevalence of category members actually engaging in such behaviors. To form broad generic beliefs based on those behaviors, it is sufficient that the relevant Ks are perceived as having a presumed common essence, which would cause them to have common dispositions. «When presented with (a) strikingly awful actions and (b) a highly essentialized group», Leslie argues, «our primitive cognitive mechanism may deliver the corresponding generalization on this basis alone» (2017: 412). Leslie (2017) supports her arguments by referring to several real-world cases. Among the most vivid and recent examples, she mentions an FBI report according to which hate crimes committed against the Muslim community in the United States would have increased by 1,600% in the months following the 9/11 attacks (Scheviz, 2002).

Importantly, Leslie (2017) claims that we do not tend to endorse striking property generalizations when they concern members of social categories in which we belong, that we perceive as socially close, or that we do not perceive as having a stereotyped shared nature.[8] Rather, in these cases, we seek more restricted generalizations involving what we perceive as better predictors, or even focus on the individual traits and behaviors of the actual perpetrators. To this respect, I believe an illuminating example is provided by a comparison between (34) and a parallel sentence about Ks that are perceived as less homogeneous, like (35):

(35) White Christians are terrorists.

[8] As Leslie notes (2017), this hypothesis is compatible with the *contact hypothesis*, according to which the interaction in personal and cooperative ways reduce inter-group prejudice (e.g., Allport, 1954; Pettigrew & Tropp, 2000).

Based on the compilation and analysis of a data set on terrorism-related incidents, Jones, Doxsee, and Harrington (2020) found that white supremacists are the most significant terror threat in the U.S.[9] Yet, people in Western countries are clearly more prone to accept social generics like (34) than social generics like (35); a white Christian engaging in terrorist acts is typically described as a "loner", as a "one hateful person", or even as "a victim of inadequate mental health resources", rather than being identified as a terrorist or based on their ethnicity and religion (e.g., Butler, 2015; Arva, Idris, & Pervez, 2017).[10] Furthermore, white Christians in the same countries are not expected to publicly disassociate themselves from the appalling actions of other white Christians. Conversely, the members of the Muslim community are subjected to such expectations, as if they were accountable for the very same actions of few individuals simply by virtue of their belonging to the same essentialized category (e.g., Butler, 2015; Dagli, 2015).

Similarly to Leslie (2017), Langton et al. (2012) and Haslanger (2011, 2014) argue that sentences like (34) might play an important role in reinforcing social prejudice and presenting social constructs as due to inherent features. In particular, since dangerous social generics might be accepted based on dispositions granted by the presumed shared nature of the category members, they might also be hard to refute. For example, a subject having prejudicial beliefs about Muslims could assert (34) and retain it true while acknowledging that almost all Muslims are not terrorists. According to the content-based theory put forward by Leslie (e.g., 2007, 2008), dangerous generic predications concern good predictors – that is, Ks need to be disposed to have the dangerous property and would display it under the right circumstances. Consequently, that subject might endorse the content of (34) despite acknowledging that the vast majority of Muslims are not terrorists. For the same reason, the acceptance of (34) would be difficult to challenge. For example, the attempt to undermine

[9] This data set contains a total number of 893 terrorist plots and attacks that occurred in the U.S. between January 1994 and May 2020 (see Jones et al., 2020).

[10] Of course, I need to point out that I am not claiming that these differences are only due to the way we speak about social categories. For example, Arab and Muslim characters are rare in movies of western production, and when included, they often play the role of villains (e.g., Semmerling, 2006; Shaheen, 2003). This certainly has a huge role in spreading prejudicial beliefs about the relevant groups in Western culture too. What I am claiming here is that, among other cultural representations, the way people speak about different social categories reflects the degree of their essentialist beliefs about those categories, and also contributes to creating and reinforcing these beliefs.

(34) by using sentences like (36) could only work if (34) is re-interpreted as a majority generic:

(36) No, most of them aren't!

However, if the proponent of (34) interpret this as a dangerous generic predication, they can always reply that although many Muslims are not terrorists, it does not mean that property is not part of their "essence"; they would exhibit that property if they were under the right circumstances because they are "disposed" to possess this property. This problem is not limited to dangerous generic predications. For example, consider the following social generic:

(37) Latin@s are lazy.

Similarly to (34), (37) might easily promote stereotyped views about the relevant Ks. If an interlocutor provides many counterexamples to combat an utterance of (37), «the speaker can then suggest that, although many Latinos aren't lazy, they tend to be» (Langton et al. 2012: 764). Moreover, if the interlocutor provides evidence that Latinos show no greater tendency toward laziness than any other group, «the speaker can then suggest that, although it is not part of the nature or essence of Latinos to be lazy, most are» (Langton et al. 2012: 764). The resistance to exceptions joined with this slide back and forth between different interpretations allows generics to resist contradiction. The difficulty of blocking generics with counterexamples of statistical significance is due to the fact that they can express characteristics, dangerous, statistical, or even normative (e.g., Leslie, 2015; Wodak, Leslie, & Rhodes, 2015; see Chapter 1, § 1.3.3) connections, but all in the same syntactic form, allowing a shift in their possible interpretation. For this reason, a subject believing a

generalization in the form of a generic will be unlikely to change their belief if contradicted and would resist possible challenges.[11]

Of course, the theoretical concerns presented in this subsection are based on the hypothesis that the content-based approach proposed by Leslie (e.g., 2007, 2008) is an accurate description of how generics are interpreted. As discussed in Chapter 1, her view as a semantic approach to generics present some theoretical difficulties. Nonetheless, her proposal received wide empirical support. While the evidence discussed in the previous section widely confirm that generics are cognitively default generalizations, other works provided evidence for the hypothesis that the interpretation of generics relies on content-based factors. Specifically, the studies carried out by Cimpian, Brandone, & Gelman (2010) confirmed that the acceptance conditions of generics are affected by the informational value of the properties they predicate. Prasada, Khemlani, Leslie, and Glucksberg (2013), instead, found that people interpret generics in a way that resembles Leslie's classification. In what follows, I briefly present these findings.

Cimpian et al. (2010) investigated the hypothesis that properties having a high informational value are more readily accepted in generic form. To this end, they examined adults' endorsement of several generic predications about novel animal kinds that were *neutral*, *dangerous*, or *distinctive*. The three item types were similar to the following, respectively:

(38)

a. Lorches have purple feathers.
   (Neutral)

b. Morseths have a silver fur that sheds particles that make it impossible to breathe.
   (Dangerous)

[11] This hypothesis is held by both Langton et al. (2012) and Haslanger (2011, 2014). However, based on her proposal concerning the power of generics to reinforce social schemes and ideologies, Haslanger (2014) further hypothesizes that the different interpretations of social categories might be affected by the presence or absence of two contextual expectations, which she calls the *Essentialist Assumption* and the *Normative Assumption*. Roughly speaking, Haslanger's idea (2014) is that by virtue of the Essentialist Assumption, the utterance of a social generic conveys that having the property F is grounded in the nature of being a K. By virtue of the Normative Assumptions, instead, the utterance of a social generics might further implicate that Ks should have F, and that they would be "defective" otherwise.

c. Reesles have distinctive blue scales that are soft, flexible, and very shiny. (Distinctive)

Participants were told what percentage of category members have a property and then asked whether the correspondent generic was true or false. For example, they were given (39a) and asked (39b):

(39)

a. 30% of lorches have purple feathers.

b. Is the following sentence true or false? Lorches have purple feathers.

Cimpian et al.'s (2010) item types were presented at the following prevalence levels: 10%, 30%, 50%, 70%, and 90%. Their results confirmed that dangerous or distinctive generic predications are more likely to be endorsed than neutral generic predications, particularly at the lowest prevalence levels. Importantly, Cimpian et al. (2010) did not observe a corresponding effect testing most- and some-quantified sentences. These findings provide further support to the view that our interpretation of generics, compared to the interpretation of quantified sentences, is sensitive to the content-based factors hypothesized by Leslie (e.g., 2007, 2008).

Another interesting confirmation of Leslie's content-based approach comes from the studies carried out by Prasada et al. (2013). In their studies, Prasada et al. (2013) found that characteristic generics, but not majority and striking property generics, involve a *principled connection* – that is, our conceptual system interpret characteristic generics as involving properties that Ks have in virtue of being Ks (see also Prasada & Dillingham, 2006, 2009). Conversely, Prasada et al. (2013) found that striking property generics involve a *causal connection* – that is, our conceptual system interprets these generics as regarding the shared nature of Ks, which cause them to be disposed to have a certain property. For example, consider the following striking property generic:

(40) Ticks carry Lyme disease.

According to Prasada et al.'s (2013) results, the interpretation of generics like (40) is related to the shared biological structure of ticks, which cause them to be disposed to carry the relevant disease. Furthermore, Prasada et al. (2013) found that majority generics involve a *statistical connection* – that is, our conceptual system interprets majority generics as regarding properties that are simply prevalent among Ks. Finally, Prasada et al. (2013) confirmed that generics do not refer to prevalence per se. For example, both striking property and majority generics were more likely to be accepted than *majority false generalizations*. As previously noted (see § 2.1.3), this category comprehends generics ascribing properties that are highly prevalent but rendered false by the presence of a salient alternative property among Ks, like (20f):

(20f) Canadians are right-handed.

However, while generics like (20f) were less likely to be accepted than both striking property and majority statistical generics, striking property generics were interpreted as referring to a lower percentage of category members than both majority statistical and false majority generics.

Both Cimpian et al. (2010) and Prasada et al. (2013) provide support to the view that the interpretation of generics reflects how our conceptual knowledge deals with properties having different informational value, and that generics correspond to different basic ways in which we represent connections between kinds and properties. In particular, they provide evidence that dangerous generic predications might be accepted based on weak evidence and concern dispositions rather than prevalence per se. In addition to providing direct evidence for Leslie's content-based approach, these findings indirectly support the proposal that social generics expressing dangerous properties might be readily endorsed and hard to challenge, worsening their already problematic impact on social cognition – that is, the promotion of social essentialism (see the previous subsection).

Given the noxious effects of social generics, several authors proposed that the use of these sentences to describe groups subjected to discrimination and prejudice should be avoided (e.g., Langton et al., 2012; Leslie 2014, 2017; Haslanger, 2011; Wodak et al., 2015). In particular, social generics might be more pernicious than quantified generalizations. Indeed, the latter do not seem to reflect our cognitively default way to express generalizations. Furthermore, generalizations having high proportion quantifiers (e.g., "most", "all") do not display the same tolerance to counterinstances of generics and can be easily undermined. Finally, essentialized categories tend to be described through generic language (e.g., Gelman et al., 2004; Rhodes et al., 2012; Segall et al., 2015), suggesting that generic beliefs might be more central to essentialist representations of kinds. In this regard, Cimpian, Gelman, and Brandone (2010) found that generics about animals are more likely to be accepted if F is an essential feature of Ks, regardless of the proportion of Ks displaying F; conversely, they observed that the acceptability of most-quantified sentences about animals depends more on considerations based on prevalence, and perhaps this discrepancy extend to the interpretation of generalizations about social categories as well.

Nevertheless, other recent proposals and empirical findings are questioning the hypothesis that social generics are especially pernicious in communication. In the next subsection, I discuss some controversies regarding the role of generics in shaping our social cognition.

2.2.4 Controversies

The proposals and empirical evidence presented thus far indicate that social generics are problematic sentences. Generics seem to facilitate the transmission of social essentialism (e.g., Rhodes et al., 2012), which foster stereotyped beliefs and prejudice (e.g., Allport, 1954); furthermore, dangerous generic predications seem to be readily accepted (e.g., Cimpian et al., 2010; Leslie, 2017) and also hard to challenge (e.g., Langton et al., 2012). For these reasons, several authors proposed we should avoid using social generics, especially when describing discriminated categories (e.g., Langton et al., 2012; Leslie, 2014, 2017; Haslanger, 2011; Wodak et al., 2015). However, it is still not clear whether social generics

are especially dangerous compared to quantified generalizations. In what follows, I consider three criticalities that have been identified concerning this hypothesis.

The first criticism is raised by Saul (2017), who argues that the results of many studies regarding the power of generics to foster social essentialism (e.g., Rhodes et al., 2012) leave it open that other constructions might promote essentialization to the same extent. Indeed, these studies frequently compare generic statements, as (41a), and specific statements, as (41b) or (41c):

(41)

a. Ks are F.

b. This K is F.

c. This is F.

According to Saul, in a context where participants are introduced to a new category, using a specific sentence conveys that having a particular property is not the norm for a member of that category. As Saul notes, it would be strange for a children book to introduce cats with sentences like (42):

(42) This cat has a tail.

Conversely, it would not be strange for the same book to describe cats with sentences like (43):

(43) This cat has black spots.

According to Saul, the reason for this discrepancy is that cats having black spots are not typical exemplars, whereas cats having a tail are. To exclude that generics lead to essentialize to a greater extent than specific sentences simply because the former have a plural form, one

might then compare the inferences elicited by social generics like (41a) with the ones elicited by plural specific sentences, like (44):

(44) These Ks are F[12].

However, this modification does not overcome the issue raised by Saul (2017). Indeed, in a context where a novel category is introduced, plural specific sentences like (44) might still suggest that the relevant property is atypical. According to Saul (2017), to support the claim that generic language plays a key role in the development of essentialist beliefs about social groups, generics should be compared to other kinds of linguistic constructions, like quantified statements signaling high proportions, as (45a) or (45b):

(45)

a. Many Ks are F.

b. Most Ks are F.

Therefore, Saul argues, more experimental data is required to conclude that generics' contribution to social essentialism is prominent. In this regard, both Cimpian and Erickson (2012) and Hoicka, Saul, Prouten, Whitehead, and Sterken (ms) actually found evidence that high proportion quantified sentences like (45a) and (45b) promote essentialist reasoning to the same extent of social generics. However, they give different explanations for this result. Cimpian and Erickson (2012) hypothesize that high proportion quantified sentences promote essentialist beliefs only indirectly. In their view, such sentences license corresponding generic beliefs about the relevant Ks. In turn, the generic beliefs so licensed would promote essentialist reasoning about the same Ks. As they argue, this hypothesis is supported by the observations that high proportion quantified sentences are often recalled as generics (e.g., Gelman et al., 2016; Leslie & Gelman, 2012; see § 2.1.3). Thus, in their view, high proportion quantified sentences do not directly lead to essentialize, but they do so by licensing a

[12] A similar solution was adopted by Rhodes et al. (2018a) in their studies concerning the role of generics in the development of social categorization (see § 2.2.3).

correspondent generic belief. Hoicka et al. (ms), instead, contend that high proportion quantified sentences could directly lead to essentialize the relevant Ks without the need to posit the formation of intermediate generic beliefs. Indeed, while high proportion quantified sentences are more likely to be recalled as generics than vice versa, they are not always misremembered as such. Consequently, if high proportion quantified sentences lead to essentialize in case they license a correspondent generic belief, they should not lead to essentialize to the same extent of generics. Future research should clarify this issue by directly controlling the extent to which the quantified sentences used to describe Ks in studies about essentialist reasoning are interpreted or recalled as generics. Together with these findings, Saul's arguments (2017), indicate that more evidence is needed to conclude whether generics are especially apt to promote essentialist beliefs, as some authors claim (e.g., Langton et al., 2012; Leslie, 2014, 2017).

The second criticism is also raised by Saul (2017; see also Ritchie, 2019). In her view, not all social generics are problematic. For example, she notices that social generics like the following might actually be used to correct discriminatory and prejudicial beliefs:

(46)

a. Women are expected to want children.
b. Black people face discrimination.
c. Gay people are subjected to violence.
d. Boys like pink too.
e. Girls play football.

As previously discussed (see § 2.2.2), it should be noted that there is evidence indicating that even generics ascribing positive features and generic corrections, which are apparently non problematic, might still convey an essentialist view about the relevant Ks (e.g., Forster-Hanson 2016, 2017). For this reason, they might paradoxically lead to accept negative stereotypes about the same groups (e.g., Kay et al., 2013) and have other detrimental effects (e.g., Cimpian et al., 2012). However, there is also evidence that social generics do not

necessarily lead to essentialize the Ks they are about. Vasilyeva and Lombrozo (2020) recently found that this side-effect of social generics can be prevented in case the association between the relevant Ks and the property F is construed in *structural* terms rather than *internalist* terms. On a structural construal, the association between being K and having F is seen as the result of stable external constraints acting on the relevant Ks in virtue of their position within the larger hierarchy of a social structure (Vasilyeva & Lombrozo, 2020; see also Haslanger, 2000). On an internalist construal, instead, the same association is seen as the result of essential features of the same Ks. To use one of the examples offered by Vasilyeva and Lombrozo (2020), consider the association between the category "immigrants" and the property "having poorly-paid jobs". A person holding a structural construal might see this association as the result of distinctive socio-economic barriers that immigrants face. A person holding an internalist construal, instead, might trace back the same association to inherent and deep presumed features of immigrants, like being less capable or not being hard working. In both cases, the association is seen as non-accidental, but it emerges for different reasons. Generics are frequently used to express non-accidental associations (e.g., Gelman, 2003), but Vasilyeva and Lombrozo (2020) hypothesized that they might permit both internalist and structural interpretation when accompanied by explanations that favor the corresponding construal. To test this hypothesis, Vasilyeva and Lombrozo (2020) presented participants with generics about a novel social category of immigrants living in a fictional country called "Kemi", like (47):

(47) Borunians hold poorly paid jobs.

Each social generic was accompanied by an explanation of why the category members have the ascribed property. For example, depending on the condition, (47) could have been explained by internalist reasons, like (48a), structural reasons, like (48b) or no reasons:

(48)

a. Borunians' inherent reluctance to take on hard and demanding jobs that may pay well, but also require over-time.

b. In order to hire a Borunian for a well-paid job, employers in Kemi are required to file complicated government paperwork.

Then, participants were asked to rate the acceptability of paraphrases of the same generics from 1 to 7. The paraphrases were internalist, like (49a), structural, like (49b), and accidental, like (49c):

(49)

a. Borunians, by virtue of being Borunians, hold poorly paid jobs.
b. Borunians, by virtue of their position in Kemi society, hold poorly paid jobs.
c. Borunians, for no particular reason, hold poorly paid jobs.

Importantly, Vasilyeva and Lombrozo (2020) found that when no reason for the association between Ks and F is provided, internalist paraphrases of generics are considered more plausible than both structural and accidental paraphrases. However, when internalist or structural reasons are provided, they promote the corresponding paraphrases of generics while suppressing the perceived acceptability of the others. Thus, while their results indicate that an internalist interpretation of generics might be the default, they also indicate that knowing that structural factors are involved in determining the association between the relevant Ks and the property F might limit the transmission of essentialist beliefs about the same Ks. In conclusion, these findings indicate that on a structural construal, social generics like the ones in (46) might actually be used without conveying an essentialist view of the relevant Ks, in case they are properly explained and disambiguated (see also Vasilyeva & Ayala-Lopez, 2019). However, while a structural construal might prevent the transmission of essentialist beliefs, it seems to me that there could be some harmful effects of social generics that would not be limited in this way. For example, consider the following pair again:

(33)

a. Boys [girls] are really good at this game.

b. There is a boy [girl] who is really good at this game.

As previously noted (see § 2.2.3), Cimpian (2010) and Cimpian et al. (2012) found that children exposed to generics like (33a) display an impaired motivation and performance than children exposed to specific sentences like (33b). Cimpian et al. (2012) hypothesize that the reason for this discrepancy is that generics like (33a) might lead to think that the relevant Ks have little control over their abilities. However, even knowing that these abilities are due to stable external constraints rather than emerging from an internal source might still discourage children from participating in the same activity. For example, one might explain to a young girl that boys are really good at a certain game because, as a consequence of a gender bias, they have been incentivized to play at it from a very young age; however, this might still lead that girl to think that, since she has never practiced it, she is unlikely to be as good as boys and might lose interest in the game. Thus, similar side-effects might be more difficult to engage than the transmission of social essentialism. Furthermore, while Vasilyeva and Lombrozo' results (2020) show an effective way to overcome the transmission of essentialist beliefs through the use of generics, they also indicate that an internalist interpretation of generics might be the default; indeed, internalist paraphrases of generics were rated as more plausible in case neither internalist nor structural explanations were provided. Furthermore, in a previous work, Vasilyeva, Gopnik, and Lombrozo (2018) found that internalist thinking is more robust than structural thinking through development. For this reasons, Vasilyeva and Lombrozo (2020) indicate as open questions the possibility that internalist thinking is less cognitively demanding than structural thinking, or that internalist explanations are perceived as more likely than structural ones. Either way, given that their basic interpretation seems to be internalist, I think that these findings support the hypothesis that generics might be particularly problematic, and that one should be very careful in using generics even on a structural construal.

Finally, a third potential criticism concerns the specific risk of dangerous generic predications. As noted in the previous subsection, the studies carried out by Cimpian et al. (2010) and Prasada et al. (2013) confirm that dangerous generic predications about animals

are more likely to be accepted based on weak evidence than other generics and that they refer to presumed shared dispositions of the categories they are about rather than to prevalence per se. Yet, to date, we do not have evidence that these findings extend to dangerous social generics. Conversely, Tasimi, Gelman, Cimpian, and Knobe (2017) found that generics ascribing threatening properties, like (50a), and generics ascribing non-threatening properties, like (50b), are accepted alike when social categories, but not artifacts or animals, are concerned:

(50)

a. Krens are dangerous.

b. Krens are helpful.

However, the non-threatening properties tested by Tasimi et al. (2017) did not have a neutral valence; in fact, they were pro-social properties, which may be interpreted differently from neutral properties, when compared with dangerous ones. Thus, it is still an open question whether social generics about dangerous properties are more easily accepted than social generics about neutral properties.

In this chapter, I discussed why investigating how we use and interpret generics constitutes a privileged way to deepen our understanding of certain aspects of our cognition. In § 2.1, I discussed the experimental data regarding the hypothesis that generics are default generalizations. In § 2.2, instead, I considered the evidence suggesting that generics facilitate the transmission of social essentialism, and that they might be especially problematic when ascribing properties that are more prominent in our conceptual knowledge. As discussed in the present subsection, some scholars argue that the empirical evidence and theoretical proposals discussed thus far does not permit to completely clarify whether social generics are more pernicious than other linguistic constructions. However, even if future research reveals that generics are not more likely to promote social essentialism than quantified generalizations, we would still have good reasons for considering the former as more pernicious and misleading than the latter. Indeed, not only generics are a common means to

describe essentialized social categories; they also present a unique combination of prevalence implications and acceptance conditions. In particular, Cimpian et al. (2010) and Brandone et al. (2015) found that generics about animals are interpreted as broadly true, yet also accepted based on weak evidence. This asymmetry is important for suggesting a mechanism by which unfounded generalizations may flourish; nonetheless, whether social generics elicit an inferential asymmetry has not been examined yet. In the next chapter, I first consider the finding of the generic asymmetry and its potential variations. Subsequently, I present my research, which concerns the robustness and generalizability of the inferential asymmetry for social generics.

3. SOCIAL GENERICS AND INFERENTIAL ASYMMETRY

In the previous chapters, I presented the semantic features and the cognitive implications of generics. An important finding lies at the crossroad between the several puzzles concerning generics. Cimpian, Brandone, & Gelman (2010) and Brandone, Gelman, & Hedglen (2015) found that generics about novel animal categories are characterized by an inferential asymmetry: they are interpreted as broadly true, yet also accepted based on weak evidence. This asymmetry is important for suggesting a mechanism by which unfounded generalizations may flourish; nonetheless, prior research has not examined whether this finding extends to the interpretation of social generics. In this chapter, I mainly discuss my work, which concerns the robustness and generalizability of the inferential asymmetry for social generics. This work was developed at the Conceptual Development Lab at the University of Michigan through collaboration with Kristan Marchak (University of Alberta), Claudia Bianchi (Vita-Salute San Raffaele University), and Susan Gelman (University of Michigan). Before focusing on our studies, I provide an overview of the finding of the inferential asymmetry and consider its potential variations as a function of domain, property type, and property valence.

3.1 Background issues

Generics have both nearly universal prevalence implications and flexible acceptance conditions. This discrepancy between the implied prevalence and truth-conditions of generics, dubbed "inferential asymmetry", was originally theorized and found by Cimpian et al. (2010). In a later work, Brandone et al. (2015) provided further support for this result, testing both adults and children. In the present section, I briefly review several key background issues regarding the hypothesis that the inferential asymmetry extends to the interpretation of social generics. In §3.1.1, I present the finding of inferential asymmetry. In

§ 3.1.2, I consider potential variations of the asymmetry as a function of domain. In § 3.1.3, I discuss potential variations of the asymmetry as a function of property type and property valence. Finally, in § 3.1.4 I briefly introduce our studies.

3.1.1 The inferential asymmetry

As discussed in the previous chapters, generics are ubiquitous and primary in our cognition. Nonetheless, the semantics of generics is not straightforward. While quantified sentences have precise truth-conditions based on the proportion of category members Ks displaying the property F, generics are not simply about quantity. For example, generics often gloss over exceptions: they express generalizations that may be judged to be true even when the vast majority of Ks lack F, unless some Ks possess a salient alternative property.

Based on the semantic features of generics, Cimpian et al. (2010) hypothesized that the interpretation of these sentences would elicit an inferential asymmetry between their prevalence implications and acceptance conditions. In the first study carried out by Cimpian et al. (2010), participants were randomly assigned to complete either an Implied prevalence task or a Truth-conditions task. In the Implied prevalence task, participants were provided with generic statements about a novel animal category and were asked to estimate what percentage of category members, from 0 to 100%, possess the ascribed property. For example, they were given (1a) and asked (2b):

(1)

a. Lorches have purple feathers.

b. What percentage of lorches do you think have purple feathers?

In the Truth-conditions task, participants were told that a certain percentage of the category members had a property and then were asked whether the corresponding generic statement was true or false. For example, they were given (2a) and asked (2b):

(2)

a. 30% of lorches have purple feathers.

b. Is the following sentence true or false? Lorches have purple feathers.

The items in this task were presented at the following prevalence levels: 10%, 30%, 50%, 70%, and 90%.

Cimpian et al. (2010) confirmed the asymmetry hypothesis: the average score in the Implied prevalence task was significantly higher than the average score in the Truth-conditions task (approximately 95% and 70%, respectively). In other words, the properties ascribed by generics were interpreted as applying to almost all members of the category, whereas the same generic statements were judged to be true even when only a small percentage of the members of the category were said to display the property. Notably, generics were distinctive in this regard; the asymmetry was not found for quantified sentences like (3), for which the average scores for truth-conditions and implied prevalence were equivalent:

(3) Most lorches have purple feathers.

An alternative explanation for the reason why participants in the first study accepted generics at low prevalence levels is that they treated these sentences as existentials – that is, sentences referring to specific members of a kind (see Chapter 1, § 1.2.2). As noted by Cimpian et al. (2010), though, this hypothesis would not account for the high percentage estimates in the generic Implied prevalence task, nor for the different acceptance levels between the types of properties. Nonetheless, Cimpian et al. (2010) carried out a second study to rule out this possibility. They tested another group of participants using a "some" Truth-conditions task, as existentials have the same truth-conditions of some-quantified sentences (see Chapter 1, § 1.2.2). Participants accepted some-quantified sentences more often than generics. This difference was significant at the 10%, 30%, and 50% levels but not at 70% and 90% (probably due, as hypothesized by Cimpian et al. (2010), to a scalar implicature

effect; see Chapter 2, § 2.1.3). The difference justified the conclusion that the BP statements in the first study were treated as generics.

In their third study, Cimpian et al. (2010) checked whether the use of accidental/temporary-sounding properties could limit the asymmetry. For example, they tested the prevalence estimations and acceptance conditions of sentences like the following:

(4)

a. Lorches have muddy feathers.

b. Ludinos have broken legs.

This study also allowed to check whether participants in the first study accepted generic sentences just as a coping strategy, given that they had to judge generics as true or false based on little information (whereas quantified statements have less ambiguous rules; see Chapter 1, § 1.2.1). Cimpian et al. (2010) found that the average implied prevalence of generics ascribing accidental/temporary-sounding properties was slightly lower than their average acceptance level. Thus, the third study confirmed that the results of the main study were not due to the absence of clear rules. Furthermore, this result added to other empirical works showing that generics do not combine effectively with accidental/temporary sounding properties (e.g., Cimpian & Markman, 2008; Gelman, 1988; Sutherland, Cimpian, Leslie, & Gelman, 2015).[1]

Cimpian et al.'s (2010) findings confirmed that generics display a distinctive combination of both nearly universal implications and flexible acceptance conditions. Recently, Brandone et al. (2015) tested whether the finding of the inferential asymmetry replicates testing

[1] For example, Cimpian and Markman (2008) observed that BP sentences referring to more generalizable properties (i.e., features that are typically shared by category members, like body parts) are more frequently interpreted as generics compared to BP sentences ascribing accidental/temporary-sounding properties. More recently, Sutherland et al. (2015) found that the use of the latter also affects the tendency to misremember quantified sentences as generics discussed in Chapter 2 (see § 2.1.3). Specifically, they found that people are more likely to recall quantified sentences as generics when they express generalizable properties than when they express accidental/temporary-sounding properties. Brandone (2017) found that evidence that we start to develop intuitions about what property types generalize broadly within the members of a category and what do not from a young age. For example, she found that even children by 6–8 years of age are less likely to extend accidental/temporary-sounding properties of an individual to other members across different domains than other types of properties (e.g., body parts, behavior, and color).

children as well as adults. They used a similar methodology to Cimpian et al. (2010), but used a more child-friendly way to test participants' intuitions about the implied prevalence and acceptance conditions of generics. For every novel animal category, they created images depicting six category members displaying the ascribed property (i.e., a physical feature) at the following ratios: 0/6 (0%), 2/6 (33%), 4/6 (66%), and 6/6 (100%). Depending on the condition they were in, participants were provided with generic statements, all-quantified statements, most-quantified statements, or some-quantified statements. Brandone et al. (2015) found that similarly to adults, children (preschoolers and 7–10-years-olds) interpret generics about animals as broadly true, yet also accept them even if only 2/6 category members display the ascribed property. However, contrary to adults, child participants in the studies carried out by Brandone et al. (2015) tended to interpret quantified sentences as if they were generics. Brandone et al. (2015) report that this finding is consistent with other studies showing that children have difficulties interpreting quantified statements even at 9 to 11 years of age (e.g., Papafragou & Schwarz, 2006; see Chapter 2, § 2.1.1). They also argue that this observation is consistent with the hypothesis that generics, unlike quantified sentences, reflect default generalizations (e.g., Gelman 2003, 2010; Leslie 2007, 2008; see Chapter 2, § 2.1.2). According to Brandone et al. (2015), children who find it difficult to interpret a quantified sentence easily embrace the generic interpretation, which is cognitively less costly. As illustrated in the previous chapter, this possibility is supported by several empirical works (e.g., Hollander, Gelman, & Star, 2002; Leslie & Gelman 2012; Meyer, Gelman, and Stilwell, 2011).

The fact that the generic asymmetry has been found with children as well as with adults suggests that it is basic in human cognition. What does elicit such asymmetry? According to Cimpian et al. (2010), the phenomenon of the inferential asymmetry might have different explanations. They first consider the possible reasons why generics have nearly universal prevalence implications. A first possibility is that generics lead people to map a certain property (e.g., plumage color) onto a certain abstract category as a whole (e.g., Lorches). However, as they note, this explanation is not sufficient to account for the high prevalence implications displayed by generics: connecting a property to a kind does not necessarily

imply that the property is widespread within the same category. Indeed, the participants in their studies also accepted generics at variable prevalence levels. Consequently, Cimpian et al. (2010) propose two more plausible explanations, which they refer to as the *default assumption* and the *pragmatic principle*, respectively. According to the default assumption, listening to a generic never heard before might automatically lead one to believe that this utterance is about a general property, as generics ascribing high prevalence features are the most prototypical. Alternatively, this phenomenon can be explained as a consequence of a pragmatic principle. Following Declerck (1991), Cimpian et al. (2010) construe this principle as the inverse of the scalar implicature effect (see Chapter 2, § 2.1.3): «If the speaker had meant to pick out a subset», as they note, «this should have been indicated by the use of another description (e.g., "some lorches"). Using an expression that is compatible with a stronger interpretation when one intends to refer to something weaker violates the norms of cooperative communication» (Cimpian et al., 2010: 1474; on the cooperative principle, see Grice, 1975). If listeners assume that the speaker is being cooperative, then the inference that the speaker intends a stronger interpretation follows. As Cimpian et al. (2010) note, the expectations driven by the default assumption and the pragmatic principle are defeasible. The relevant inferences might be blocked by the use of certain semantic and pragmatic cues (e.g., the use of accidental/temporary sounding properties; see also Chapter 2, § 2.1.2).

According to Cimpian et al. (2010), the reason why generics have also flexible acceptance conditions is less related to the reconstruction of the communicative intentions of the speaker. In their view, to accept or not a given generic, a different two-step process is required. The listeners have to determine a) the range of circumstances in which that generic would be true and then b) whether the evidence available is included in that range. Since generics lack explicit quantifiers, the range of circumstances in which these sentences would be true will be often ambiguous. Such ambiguity may result in a widening of the acceptable range of circumstances, making even relatively low evidence likely to fall under that range. Furthermore, as also highlighted by Brandone et al. (2015), since generics are not simply about the proportion of Ks displaying F, their evaluation incorporates considerations linked to our conceptual knowledge about categories that, under certain circumstances,

counterbalances the effect of low prevalence (see the next section, § 3.2.2; Chapter 1, § 1.3.3; Chapter 2, § 2.2.3). For example, in the representation of an animal category, a property like "having purple feathers" might be more prominent than "having muddy feathers"; for this reason, a generic expressing the former might be more likely to be accepted even based on weak evidence than a generic expressing the latter.

Cimpian et al. (2010) proposed that the generic asymmetry may be relevant to several real-world phenomena, such as knowledge acquisition through generic language, or propaganda in political discourse. The generic asymmetry also raises concerns regarding the use of social generics, which are a common means to express stereotypes (Gelman, Taylor, & Nguyen, 2004). If social generics elicit an inferential asymmetry, they could be especially pernicious and misleading, as a property that is true only of a small number of a group of people may be assumed to be broadly representative (e.g., Bian & Cimpian, 2017; Brandone et al. 2015; Cimpian et al., 2010; Hammond & Cimpian, 2017; Langton, Haslanger, & Anderson, 2012; Rosola & Cella, 2020; see § 3.6.4). However, the generic asymmetry has not yet been examined in social categories. In the next subsection, I consider some issues concerning whether the inferential asymmetry elicited by generics differs as a function of domain.

3.1.2 The role of domain in the interpretation of generics

Among the most far-reaching implications of the asymmetry is its potential for helping to explain the transmission and acceptance of unwarranted generalizations, such as stereotypes about social categories. However, as noted in the previous subsection, to date nothing is known about whether the asymmetry would be found in this domain.

When examining generics about social categories, we might expect to see the same inferential asymmetry as was observed for animal kinds, given that social categories often are treated as if they were natural animal kinds (Prentice & Miller, 2007; Rhodes & Gelman, 2009; Rothbart & Taylor, 1992). Consistent with this possibility, generic language promotes essentialist reasoning in both domains (Foster-Hanson, Leslie, & Rhodes, 2016, 2019;

Gelman, Ware, & Kleinberg, 2010; Leshin, Leslie, & Rhodes, 2020; Rhodes, Leslie, Saunders, Dunham, & Cimpian, 2018b; Rhodes, Leslie, & Tworek, 2012; see Chapter 2, § 2.2).

A second possibility is that we may find a greater asymmetry in the social domain, if people are especially prone to endorse generic statements about social categories on the basis of minimal evidence. Specifically, we may expect people to judge social generics to be true at lower prevalence levels than they judge generics about animals to be true, resulting (all other things being equal) in a greater asymmetry for social generics. This seems plausible, given that stereotypes operate by generalizing well beyond the evidence (e.g., Allport, 1954; Wodak, Leslie, & Rhodes, 2015). For example, people endorse as true the stereotype that "boys don't cry", even though most boys do in fact cry (Leslie, 2015; Wodak et al., 2015).[2] However, a third possibility is that we may instead expect to see less asymmetry for social generics, or even no asymmetry at all, due to differences in the structure of social and animal categories. In particular, we may find lower implied prevalence estimates for social generics than generics about animals. In contrast to animal kinds, whose members are often highly similar to one another (e.g., different skunks are highly similar in appearance and behavior), members of a social kind may be more variable (e.g., Brandone, 2017; Nisbett, Krantz, Jepson, & Kunda, 1983). For example, girls may differ from one another in age, race, ethnicity, preferences, dwellings, clothing, dietary preferences, abilities, languages spoken, etc. If people expect more variability among members of a social category, then generics regarding social categories may elicit lower implied prevalence ratings, and consequently lead to less of an asymmetry.

In short, then, it is an open question whether the inferential asymmetry holds for social generics. In the next subsection, I discuss potential variations of the asymmetry as a function of property type and property valence.

[2] In particular, as discussed in Chapter 1 (§ 1.3.3), Leslie (2015) points out that certain social generics, like "Boys don't cry", can get a normative reading. Under this interpretation, they do not concern actual individuals, but ideal ones. Hence, if a listener interprets "Boys don't cry" as expressing a norm, they will not rely on instances of crying boys to evaluate this sentence.

3.1.3 The role of property type and valence in the interpretation of generics

The generic asymmetry might also vary depending on the type of properties expressed. Prior research has documented that people have domain-specific beliefs about the homogeneity of properties with respect to categories (e.g., Brandone, 2017; Gelman, 1988; Nisbett et al., 1983). For example, Nisbett et al. (1983) found that for social categories, people were more confident when generalizing skin color than body weight, based on a few members of the category. The extent to which they generalized skin color for a social category was also greater than the extent to which they generalized the color and the nesting habit of a novel bird category. These differences may affect the interpretation of generics about animal and social categories with respect to the ascribed properties. In particular, past research on the asymmetry was limited to physical properties of animal kinds, which may have uniquely broad inferential potential compared to non-physical properties, such as behaviors. For example, all birds have feathers [physical property], but not all birds fly (e.g., penguins and chickens) [behavioral property]. Thus, the inferential asymmetry found by Cimpian et al. (2010) and Brandone et al. (2015) for generics ascribing physical properties to animal categories may not necessarily hold for generics ascribing non-physical properties or, more generally, for social generics ascribing physical and non-physical properties.

Another important issue involves the informational value of the properties being expressed. As previously discussed (see Chapter 1, § 1.3.3; Chapter 2, § 2.2.3), Leslie (e.g., 2008) hypothesized that generics ascribing *distinctive* and *dangerous* properties are more easily accepted than generics ascribing *neutral* properties. Properties that are perceived as distinctive or dangerous have relatively high informational value, and, in turn, might be more prominent in our conceptual knowledge, and so more readily accepted in generic form. As noted in the previous chapter (see § 2.2.3), to test this hypothesis, Cimpian et al. (2010) examined people's endorsement of properties of novel animal kinds that were neutral (e.g., "have purple feathers"), dangerous (e.g., "have a silver fur that sheds particles that make it impossible to breathe"), or distinctive (e.g., "have distinctive blue scales that are soft, flexible, and very shiny"). Importantly, they found that generics expressing dangerous or distinctive properties of animals were more likely to be judged true than generics expressing

neutral properties. It is unclear, however, whether this tendency would generalize to social categories. Tasimi, Gelman, Cimpian, and Knobe (2017) found that generics ascribing threatening properties (e.g., "are dangerous") and generics ascribing non-threatening properties (e.g., "are helpful") were accepted alike when social categories, but not artifact or animal categories, were concerned. However, as also noted in the previous chapter (see § 2.2.4), the non-threatening properties tested by Tasimi et al. (2017) were in fact pro-social properties; pro-social properties may be interpreted differently from neutral properties, when compared with dangerous ones. Consequently, whether social generics expressing dangerous properties are more easily accepted than social generics expressing neutral properties requires further examination.

Finally, a related issue concerns people's willingness to generalize a property based on whether the generic predication is dangerous or neutral. As previously discussed (see Chapter 2, § 2.2.3), Prasada, Khemlani, Leslie, and Glucksberg (2013) found that dangerous generic predications about familiar natural and artifact categories were interpreted as referring to common *dispositions* rather than to prevalence per se. For example, the interpretation of a generic like (5),

(5) Ticks carry Lyme disease.

is related to the shared biological structure of ticks, which cause them to be disposed to carry the relevant disease. Such interpretation does not take into account how prevalent the ascribed property is; indeed, the actual proportion of category members displaying the dangerous property depends on whether determined enabling conditions are present (e.g., whether ticks feed on infected animals), which may be rare or even absent. Such statements led to lower prevalence estimates than other generic predications. Similarly, Lazaridou-Chatzigoga, Katsos, and Stockall (2019) found that both adults and children were more willing to extend the properties ascribed by neutral generic predications about novel creatures, like (6a), to a new member of the category than the properties ascribed by dangerous generic predications, like (6b):

(6)

a. Ackles love to play with toys.

b. Ackles love to play with fire.

However, as Lazaridou-Chatzigoga et al. (2019) noted, the properties they tested were more child-friendly than those of Cimpian et al. (2010) and thus potentially less salient.[3] Moreover, many of their dangerous properties were likely to be perceived as dangerous for the category members only. It also may be that baseline assumptions played a role (e.g., in general, playing with toys is considered more enjoyable than playing with fire). In conclusion, another open question is whether the prevalence estimations elicited by dangerous generic predications about novel animals or people differ from those based on neutral generic predications.

3.1.4 The present studies

Together with Marchak, Bianchi, and Gelman, I conducted four preregistered studies to test whether the inferential asymmetry found by Cimpian et al. (2010) with generics about animals extends to social generics. Furthermore, we examined whether different property types (physical vs. non-physical) and valence (neutral vs. dangerous) affect people's judgments. In Study 1, we investigated whether the finding of the inferential asymmetry obtained with generics about animal categories replicates using an improved methodology and a larger online sample. In Study 2, we tested whether the generic asymmetry varies as a function of domain (animals vs. people) and property type (physical vs. non-physical). In Studies 3 and 3a, we tested whether the generic asymmetry varies as a function of domain (animals vs. people) and property valence (neutral vs. dangerous).

[3] The full list of the neutral properties tested by Lazaridou-Chatzigoga et al. (2019) is "love to play with toys", "love to talk to their mothers", "love to run through parks", "love to draw their names", "love to feel safe", "love to make new games", "love to play with cats", "love to make people sing". The dangerous properties they tested, instead, are "love to play with fire", "love to scare their mothers", "love to smash through walls", "love to shout their names", "love to feel afraid", "love to cheat at games", "love to play with snakes", "love to make people angry".

3.2 Study 1

In Study 1, we examined whether the finding of the asymmetry replicates with a larger online sample. Furthermore, we examined whether this finding holds when a 100% prevalence level was included in the Truth-Conditions task, as the prevalence levels used by Cimpian et al. (2010) in their Truth-conditions task did not include 100%. The inclusion of a 100% level in our study simplified the interpretation of participants' responses, because we did not need to impute a score of 100% to participants that said "no" to all questions, as was done in Cimpian et al. (2010). The preregistration of this study is available at https://aspredicted.org/blind.php?x=mw5pn8.

3.2.1 Method

3.2.1.1 Participants

One-hundred and twelve adults from the United States (60 men, 52 women; Mean age = 42.99 years; range = 24–75 years) completed the study online via Amazon Mechanical Turk (MTurk) for $0.40.[4] Participants were randomly assigned to either the *Implied Prevalence* task (n = 55) or the *Truth-Conditions* task (n = 57). In this and each subsequent study, participants had been granted Master Worker status by MTurk, had a US IP address, 1000+ approved HITs, and a 99%+ HIT approval rate. Participants were 83% White, 8% Asian or Asian American, 5% Black or African American, 2% Latino or Hispanic, and 2% Multiracial/Multiethnic.[5] Five additional participants were tested and excluded from the final sample for having non-US IP addresses. One other participant was tested and excluded for having a duplicate IP address.

[4] In the pre-registration of subsequent studies, we reported that 175 participants took part in Study 1. There are two reasons for the discrepancy between the number of participants reported here and in the pre-registrations. First, 59 participants completed a Truth-Conditions task that matched the original design of Cimpian et al. (2010) (i.e., without a 100% prevalence level). I report the results of that task in the Appendix (see "Study 1"). Second, we initially believed that only one participant in this study had a non-US IP address. However, we realized after completing the pre-registration that four additional participants in Study 1 had a non-US IP address.

[5] In this and subsequent studies, participants were asked to self-identify their race/ethnicity. The response options were a) American Indian or Alaska Native, b) Asian or Asian American, c) Black or African American, d) Latino or Hispanic, e) Middle Eastern or North African, f) Native Hawaiian or Other Pacific Islander, e) White, and/or f) Not listed.

3.2.1.2 Materials and procedure

In all studies, novel categories and novel labels were used, to ensure that participants were not simply retrieving learned facts. We created a list of 12 items, each consisting of a novel label for an animal category and a property that described the color of a body part of the animal (e.g., "MORSETHS have silver fur"). The list consisted of the 10 plain properties from Cimpian et al. (2010) and two additional items that we created for this study. (This was done to include two items at the 100% prevalence level; for a complete list of items, see Appendix S1). Also, in contrast to Cimpian et al. (2010), assignment of label to property was randomized for each participant. Moreover, in contrast to Cimpian et al., the labels for the novel categories in our studies were presented entirely in uppercase letters. This permitted consistent presentation for animals and people in subsequent studies, as otherwise category labels for animals typically begin with a lowercase letter (e.g., dogs), whereas category labels for people typically begin with a capital letter (e.g., Canadians). For an example of an item as it would appear to participants, see Table 1.

Implied Prevalence task	*Truth-Conditions task*
Information: MORSETHS have silver fur. Question: What percentage of MORSETHS have silver fur?	Information: 30% of MORSETHS have silver fur. Question: Is the following sentence true or false? **MORSETHS have silver fur.**

Table 1. Sample item from Study 1 in the *Implied Prevalence* and *Truth-Conditions* tasks.

At the beginning of the survey, each participant read the same introductory text as presented in Cimpian et al. (2010):

> "In this study, we will tell you about some animals that live on a remote island. This island is very large and has many different animals on it. For each item, you will be

given some information and asked a question. Please try to answer our questions to the best of your ability."

In the *Implied Prevalence* task, participants were presented with a generic statement and then asked to estimate the implied prevalence of the property described in the statement, from 0 to 100%. Participants in this task received 10 items randomly selected from the list of 12 items. (We provided 10 items instead of 12 in order to equate the items with that of the original Truth-Conditions task mentioned in footnote 4 and presented in the Appendix; see "Study 1".) In the *Truth-Conditions* task, participants judged whether generic statements were "true" or "false" based on the prevalence level of an ascribed property, with each of the following prevalence levels presented twice: 10%, 30%, 50%, 70%, 90%, 100%. In this task, participants received all 12 items. The order of the items in both tasks was randomized for each participant.

3.2.2 Results and discussion

3.2.2.1 Data coding

To enable comparison of participants' responses in the *Implied Prevalence* task to that in the *Truth-Conditions* task, in this and subsequent studies, we calculated mean prevalence scores for participants in each task (based on the coding scheme used in Cimpian et al., 2010). In the *Implied Prevalence* task, we averaged participants' responses out of 100% across the 10 items. In the *Truth-Conditions* task, we converted participants' "true"/"false" judgments to a mean prevalence level that led to "true" responses. To calculate this score, we added the percentage level of items that a participant judged to be "true" and divided this score by the total number of items that the participant judged to be "true". For example, a participant's score would be 80% if they selected "true" on two items, one where the prevalence was 70% and the other where the prevalence level was 90%.

3.2.2.2 Did the generic asymmetry replicate with a larger online sample and including the 100% level in the Truth-Conditions task?

We conducted an independent sample *t*-test of the mean prevalence ratings in the two tasks. As in Cimpian et al. (2010), we observed an asymmetry between ratings, with participants providing higher mean ratings in the *Implied Prevalence* task ($M = 95.6\%$) than in the *Truth-Conditions* task ($M = 69.5\%$), $t(110) = 10.49$, $p < .001$, $d = 1.98$ (see Figure 1).

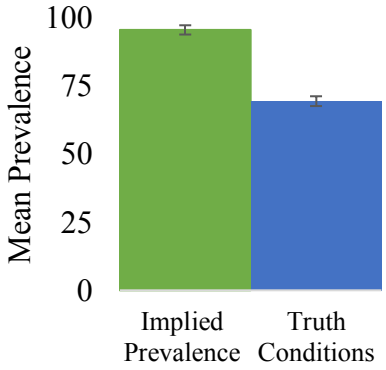


Figure 1. Study 1, the mean prevalence implied by generic statements (*Implied Prevalence* task, green bar) and the mean prevalence that led participants to accept the same generic statements (*Truth-Conditions* task, blue bar). Error bars represent one standard error.

3.2.3 Conclusions

The results of Study 1 provide clear evidence of a generic asymmetry about physical features of animal categories when the methods of Cimpian et al. (2010) were replicated with a larger sample. In the Implied Prevalence task, participants interpreted generics as referring to nearly all members of the kind. In the Truth-Conditions task, instead, they judged these same generics to be true at substantially lower prevalence levels. These results suggest that the generic asymmetry is robust, even with an improved methodology that includes a 100% level in the Truth-Conditions task.

3.3 Study 2

In Study 2, we investigated whether a generic asymmetry would be found for social generics, as well as for generics about animal categories. Since previous works on the asymmetry were limited to physical properties of animal categories, which may have different inductive power than non-physical properties in different domains (e.g., Brandone, 2017; see § 3.1.3), we also examined how property type (physical vs. non-physical) affected the pattern of responses in social and animal categories. In the Truth-Conditions task, we also specifically examined whether domain and property type affected judgments of generics' acceptability at different prevalence levels. The preregistration of this study is available at https://aspredicted.org/blind.php?x=ru2pa8.

3.3.1 Method

3.3.1.1 Participants

Two-hundred and thirteen adults from the United States (96 men, 116 women, 1 undisclosed; Mean age = 40.72 years; range = 21–68 years) completed the study on MTurk and were paid $1. Participants were randomly assigned to one of four conditions: *Animals–Implied Prevalence* (*n* = 52), *People–Implied Prevalence* (*n* = 53), *Animals–Truth-Conditions* (*n* = 56), *People–Truth-Conditions* (*n* = 52). Participants were 75% White, 8% Asian or Asian American, 7% Black or African American, 6% Multiracial/Multiethnic, 3% Latino or Hispanic, 0.5% Ashkenazi Jew, 0.5% not listed/other, and 0.5% undisclosed. Twenty-one additional participants were tested and excluded because they failed the manipulation check (*n* = 10 in the *People–Implied Prevalence* task, and *n* = 11 in the *People–Truth-Conditions* task; see § 3.3.1.2). Three participants were excluded from the final sample for having non-US IP addresses. Two other participants were tested and excluded from the final sample since

they had already participated.[6] One participant was also tested and excluded from the final sample for not being a native speaker of English.

3.3.1.2 Materials and procedure

We created a list of 24 items: 12 described a physical property of the category (e.g., "Ks have large tonsils"), and 12 described a non-physical property (e.g., "Ks sleep under trees"). For each participant, the properties were randomly assigned to one of 24 labels (see Appendix S2 for a complete list of the labels and items used in this study).

We adapted the introductory text from Cimpian et al. (2010) and Study 1 in the following way:

> "In this study, we will tell you about some **[animals]**/**[people]** that live on a remote island. This island is very large, and has many different **[animals]**/**[people]** on it. It is roughly the size of Alaska, and has a lot of geographical, climatic, and environmental variety. For each item, you will be given some information and asked a question. Please try to answer our questions to the best of your ability. Remember: the following questions are about **[animals]**/**[people]** on an island."

We described the island as being "roughly the size of Alaska" and having "a lot of geographical, climatic, and environmental variety" to avoid the possibility that participants would imagine a stereotypical, small tropical island. We were specifically concerned about the influence of beliefs about the size of the island on responses in the *People* conditions because (1) participants may already have pre-existing beliefs about groups of people living on tropical islands, and (2) they may also reject the premise that numerous different groups of people could live on a small island.

[6] One participant had completed the norming study that we conducted to select the items for Study 3. (The norming study was the first study that we conducted.) The other participant had already participated in Study 1.

As in Study 1, participants completed either the *Implied Prevalence* task or the *Truth-Conditions* task. In both tasks, participants were presented with all 24 items, but the order of the items was randomized for each participant. In the *Truth-Conditions* task, participants judged two physical properties and two non-physical properties at each of the following prevalence levels: 10%, 30%, 50%, 70%, 90%, 100%.

After each task, we also included a manipulation check to determine whether participants remembered whether the questions were about animals or people (depending on the condition they were in). We reminded them that they would see a sentence that they had originally seen at the beginning of the study, and asked them to fill in the blank in the following sentence:

"Remember: the following questions are about _____ on an island."

Participants who failed this manipulation check were excluded from the final sample (see § 3.3.1.1).

At the end of the study, we asked participants to complete an abbreviated 13-item social desirability measure (Reynolds, 1982) to examine whether responses on this measure were predictive of people's responses in the *People* conditions in the main task.

3.3.2 Results and discussion

3.3.2.2 Did the generic asymmetry vary as a function of domain and/or property type?

To explore this question, we conducted a repeated measures analysis of variance (ANOVA) on the mean prevalence ratings, with task (Implied Prevalence vs. Truth-Conditions) and domain (animals vs. people) as between-subjects factors and property type (physical vs. non-physical) as a within-subjects factor (see Figure 2). We observed a main effect of task, with ratings overall higher in the *Implied Prevalence* task ($M = 95.0\%$) than the *Truth-Conditions* task ($M = 74.8\%$), $F(1, 209) = 141.96$, $p < .001$, $\eta_p^2 = .40$. No other main effects or interactions were significant.

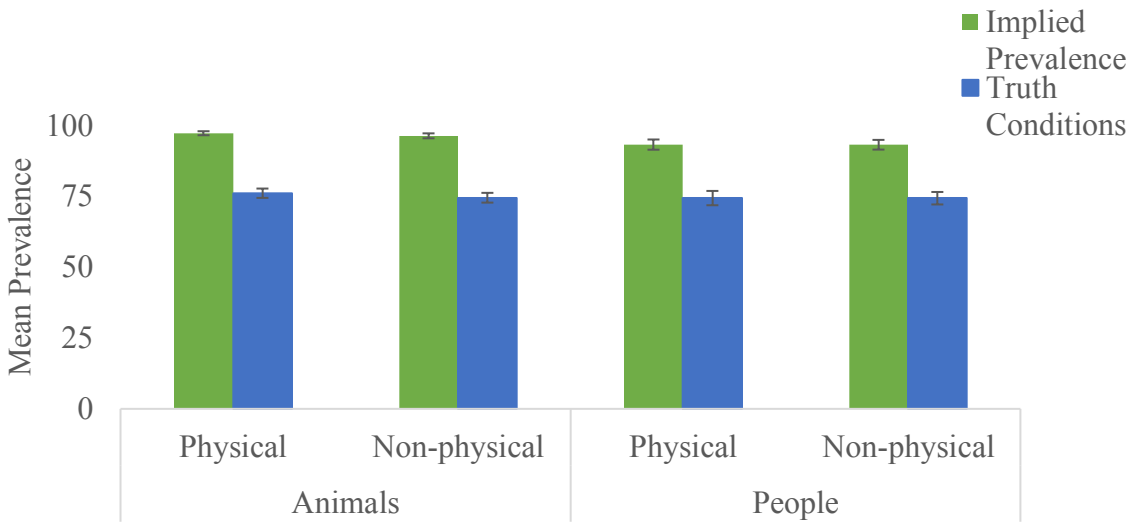


Figure 2. Study 2, the mean prevalence implied by generic statements (*Implied Prevalence* task, green bars) and the mean prevalence that led participants to accept the same generic statements (*Truth-Conditions* task, blue bars). Error bars represent one standard error.

We did not observe a significant correlation between a person's score on the social desirability scale and their mean prevalence score in either the *People–Implied Prevalence* or *People–Truth-Conditions* tasks for physical or non-physical properties, all *p*s > .05.

3.3.2.3 Did domain and property type affect generics' acceptability at different prevalence levels?

We submitted participants' "true/false" responses in the Truth-Conditions task to a logistic mixed-effects model using the glmer command in the lme4 package in R (Bates, 2007). This analysis was not preregistered, but was included in this and subsequent studies in order to compare our results with those of Cimpian et al. (2010). In this model, we included domain (animals = 0; people = 1; between subject), property type (non-physical = 0; physical = 1; within subject), and prevalence (.1, .3, .5, .7, .9, 1; within subject) and their interactions as predictors (see Table 2). All predictors were mean centered. We also included *participant* as

a random effect.[7] We observed a main effect of prevalence, indicating that generic sentences were more likely to be judged to be true for higher than lower prevalence levels. We additionally observed an interaction between domain and prevalence.

Fixed Effects	**Estimate**	***SE***	***p*-value**
(Intercept)	2.30	0.39	< .001
Domain	0.09	0.75	.90
Property Type	-0.25	0.15	.11
Prevalence	9.26	0.45	< .001
Domain x Property Type	-0.08	0.31	.79
Domain x Prevalence	-2.29	0.82	.01
Property Type x Prevalence	-0.03	0.53	.95
Domain x Property Type x Prevalence	-1.66	1.05	.12
Random Effect		**SD**	
Participant	Intercept	3.68	

Table 2. Logistic regression predicting "true"/ "false" judgments, based on domain, property type, and prevalence and their interactions in Study 2.

To explore the interaction (see Figure 3), we conducted post-hoc tests that revealed that participants were numerically, but not significantly, more likely to endorse social generics than generics about animals at the 10% level (Average Marginal Effect (AME) = 0.07, *SE* = 0.05, *p* = .12, 95% CI = -0.02, 0.17), 30% level (AME = 0.08, *SE* = 0.08, *p* = .32, 95% CI = -0.07, 0.22), and at the 50% level, (AME = 0.03, *SE* = 0.08, *p* = .71, 95% CI = -0.13, 0.19). In contrast, we observed that participants were numerically, but not significantly, more likely to endorse generics about animals than social generics at higher prevalence levels: 70% level (AME = -0.01, *SE* = 0.05, *p* = .81, 95% CI = -0.12, 0.09), 90% level (AME = -0.03, *SE* = 0.04, *p* = .45, 95% CI = -0.12, 0.05), and 100% level (AME = -0.04, *SE* = 0.04, *p* = .32, 95% CI = -0.11, 0.04). For a complete overview of the mean endorsement percentages at each

[7] We additionally fit a model including item as a random effect; however, we found that the estimate for the SD of the intercept for item was zero, so we omitted item as a random effect in the final model in this and subsequent studies.

prevalence level by domain and property type, see "Truth-Conditions – Study 2" in the Appendix.

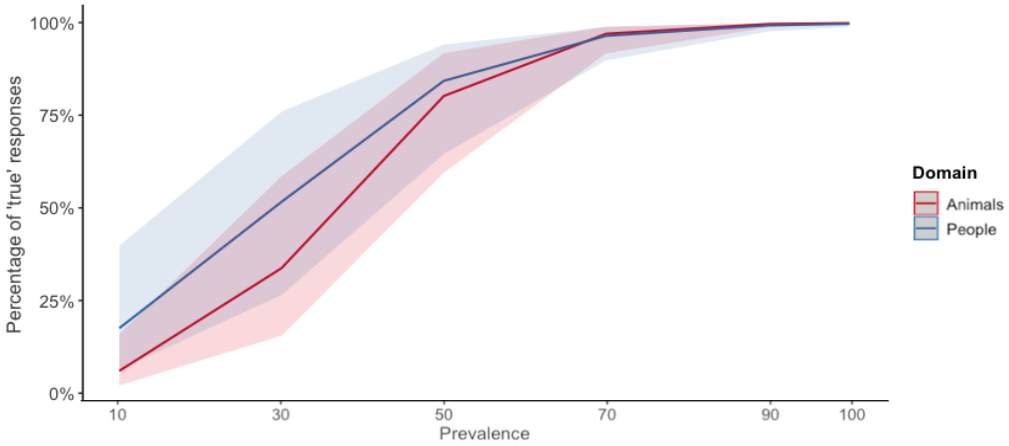


Figure 3. Mean percentage of "true" responses in Truth-Conditions task in Study 2 by domain and prevalence level.

3.3.3 Conclusions

Overall, the results of Study 2 support the robustness and generalizability of the generic asymmetry. In particular, these results show that the inferential asymmetry elicited by social generics is not significantly different from the one elicited by generics about animals. Furthermore, we observed that the two property types (physical and non-physical) did not elicit different degrees of asymmetry. Finally, despite the overall lack of domain or property type effects, the acceptability of generics in the Truth-Conditions task was affected by a small but measurable interaction between domain and prevalence: social generics were more likely than generics about animals to be accepted at the lowest prevalence levels, whereas generics about animals were more likely than social generics to be accepted at the highest prevalence levels. The reasons why we observed this effect are unclear, as we did not observe any effect related to domain in the Implied Prevalence task. Thus, this result should be replicated before it is ascribed much significance.

The reason for the overall lack of domain or property type effects in this study is also unclear. A first possibility is that participants perceived the various category/property pairings as having equivalent within-category homogeneity. In a supplementary study, we explored this possibility by asking participants whether the same properties presented in Study 2 could generalize from three instances of a category to other members of the same category (see "Study 2" in the Appendix). Participants judged the properties in our study to be more homogeneous in animal than social categories (e.g., they judged members of an animal category to be more alike in terms of their color or habits than members of a social category). In contrast, we observed no difference between estimates for physical and non-physical properties. These results suggests that the null property type effect may be due to the properties having equivalent and high within-category homogeneity. However, it does not explain the null domain effect in this study, since we observed that properties in our studies were judged more homogenous in animal than people categories.

Alternatively, the null property type effect might be attributed to the physical and non-physical properties being perceived as having a similar informational value. This possibility is consistent with Leslie's (2008) hypothesis and Cimpian et al.'s (2010) finding that generics are evaluated differently according to the informational values of the properties they ascribe. Accordingly, when two different property types are perceived as having similar informational values, they are evaluated alike in generic form. While this hypothesis might account for the lack of a property type effect, it does not account for the null domain effect in this study. Since we observed a domain effect in the supplementary study, we speculate that shifting from specific to generic claims made domain differences disappear.

Previous research indicates that dangerous information is featured as more prominent than neutral information in our kind representations (e.g., Cimpian et al., 2010; Leslie, 2008). For this reason, in Study 3, we provided a stronger test of the role of informational value by varying property valence, and tested if/how this factor affects the generic asymmetry and the pattern of responses in the two domains.

3.4 Study 3

In Study 3, we again tested for a generic asymmetry in two category domains (people and animals), but this time we also varied property valence (dangerous vs. neutral). In the *Truth-Conditions* task, we also examined whether domain and property valence affected generics' acceptability at different prevalence levels. We used a similar methodology to Study 2. The preregistration of this study is available at https://aspredicted.org/blind.php?x=zr3sy3.

Cimpian et al. (2010) originally examined whether the inferential asymmetry and the acceptance conditions of generics about animals vary as a function of property valence. In addition to the inclusion of social generics, our methodology differs from that of Cimpian et al. (2010) in a number of important respects.

First, in Cimpian et al.'s (2010) studies, the interpretation of neutral generic predications was compared to the interpretation of dangerous generic predications that were either a) both dangerous *and* distinctive (e.g., "Reesles have blue scales that secrete a strong venom that kills you on the spot. No other animals have this kind of scales") or b) intermingled with generics that varied in distinctiveness, potentially priming participants to think about this particular factor throughout. For this reason, in the present study we tested only dangerous and neutral properties. Our properties were also pretested to validate their status as dangerous or neutral.

Second, in Cimpian et al.'s (2010) studies, the neutral generic statements were substantially shorter than the other types of generic statements. For this reason, in their first study Cimpian et al. (2010) additionally included non-distinctive control generic sentences of approximately the same length as the dangerous/distinctive sentences to examine whether simply providing more information could affect participants' evaluations. Cimpian et al. (2010) observed that dangerous/distinctive generics were more likely to be accepted than both neutral and non-distinctive generics. For this reason, they concluded that participants' responses were due to the informational value of items rather than the items' length. However, this represented only an indirect control of the impact of item length; indeed, Cimpian et al. (2010) needed to infer from the observation that dangerous/distinctive generics

were more likely to be accepted than both neutral and non-distinctive generics that the length difference between dangerous/distinctive generics and neutral generics did not explain their findings. In the present study, we tested dangerous and neutral generic sentences that were roughly equivalent in length, with predicates ranging from 2 to 6 words to provide a more direct control of item length.

Third, we took care to ensure that neither dangerous nor neutral items in the present study included generic-you, given that this expression conveys norms and broad generalizations (Orvell, Kross, & Gelman, 2017, 2019, 2020). This is in contrast to Cimpian et al. (2010), which included generic-you for dangerous but not neutral generic statements (e.g., "have red scales that secrete a strong venom that kills *you* on the spot"; italics added).

3.4.1 Method

3.4.1.1 Participants

Two-hundred and twenty-nine adults from the United States (93 men, 133 women, 2 undisclosed, and 1 gender-fluid; Mean age = 41.00 years; range = 24–72 years) completed the study online and were paid $1. Participants were randomly assigned to one of four conditions: *Animals–Implied Prevalence* (n = 54), *People–Implied Prevalence* (n = 66), *Animals–Truth-Conditions* (n = 57), or *People–Truth-Conditions* (n = 52). Participants were 81% White, 6% Black or African American, 5% Multiethnic, 5% Asian or Asian American, 2% Latino or Hispanic, 0.4% Native Hawaiian or Other Pacific Islander, and 0.4% not listed. Nineteen other participants were tested and excluded from the final sample because they failed the manipulation check (n = 2 in the *Animals–Implied Prevalence*, n = 2 in the *People–Implied Prevalence, n* = 3 in the *Animals–Truth-Conditions*, and n = 12 in the *People–Truth-Conditions*, tasks). Four participants were tested and excluded from the final sample for having duplicate IP addresses. Three participants were tested and excluded from the final sample because they were not native speakers of English. One additional participant was tested and excluded for having a non-US IP address.

3.4.1.2 Materials and procedure

An additional 97 participants completed a norming study to allow us to select dangerous and neutral items. We presented participants with the same introductory text as in Study 2. Half of the participants (n = 50) were asked about animal categories and half (n = 47) were asked about people categories.[8] Participants were presented with 60 items: twenty intended to be dangerous (e.g., "Ks hunt strangers"), twenty intended to be neutral (e.g., "Ks sleep under trees"), and twenty intended to be safe (e.g., "Ks help others"). They were then asked to rate how dangerous the category is (e.g., "How dangerous are Ks?") on a scale of 1 (safe) to 7 (dangerous), with the midpoint (4) as neutral. Each item consisted of a novel label paired with one of the 60 properties; furthermore, the order of these items was randomized for each participant. Based on the norming study, we selected 24 items for use in the main experiment: 12 neutral items (e.g., "Ks hide underground"; M rating = 3.58, SD = .75) and 12 dangerous items (e.g., "Ks carry a deadly virus"; M rating = 6.71, SD = .47), t (96) = 33.18, $p < .001$, d = 3.37; see Appendix S3 for a complete list of the items. For both the dangerous and neutral properties, half of the items included descriptions of traits (e.g., "are messy", "are dangerous") and the other half included descriptions of behaviors (e.g., "sleep under trees", "hunt strangers"). We included these factors because we were interested in examining whether the interpretation of traits differs from the interpretation of behaviors. A trait might be manifested by a great variety of properties, whereas a behavior corresponds to more specific instantiations of properties. For this reason, traits might elicit more varied beliefs about the category members than behaviors, affecting participants' responses.

Unlike the dangerous properties used by Cimpian et al. (2010), our dangerous properties had approximately the same number of words as the neutral properties. For each participant, the properties were randomly associated with one of 24 labels from Study 2 (see Appendix S2 for a list of the labels). The order of the items was randomized for each participant.

Participants were presented with the same introductory text as in Study 2 and then completed either the *Implied Prevalence* task or the *Truth-Conditions* task. After the task,

[8] In initial analyses, we explored whether there were domain differences in ratings. However, we observed no significant impact of domain; thus, we collapsed across this factor in subsequent analyses of data from the norming study.

participants were also asked to complete the same manipulation check as in Study 2 and Reynolds' (1982) social desirability measure.

3.4.2 Results and discussion

3.4.2.2 Did the generic asymmetry vary as a function of domain and/or property valence?

We conducted a repeated measures ANOVA focused on the mean prevalence ratings, with task (Implied Prevalence vs. Truth-Conditions) and domain (animals vs. people) as between-subjects factors, and property valence (dangerous vs. neutral) as a within-subjects factor (see Figure 4). We found a main effect of task, with ratings overall higher in the *Implied Prevalence* task ($M = 88.8\%$) than the *Truth-Conditions* task ($M = 74.2\%$), $F(1, 225) = 51.07$, $p < .001$, $\eta_p^2 = .19$. We additionally observed a main effect of property valence, $F(1, 225) = 27.30$, $p < .001$, $\eta_p^2 = .11$, with higher ratings for neutral properties ($M = 83.5\%$) than dangerous properties ($M = 79.5\%$). However, these main effects need to be interpreted within the context of an interaction between task and property valence, $F(1, 225) = 23.41$, $p < .001$, $\eta_p^2 = .09$.

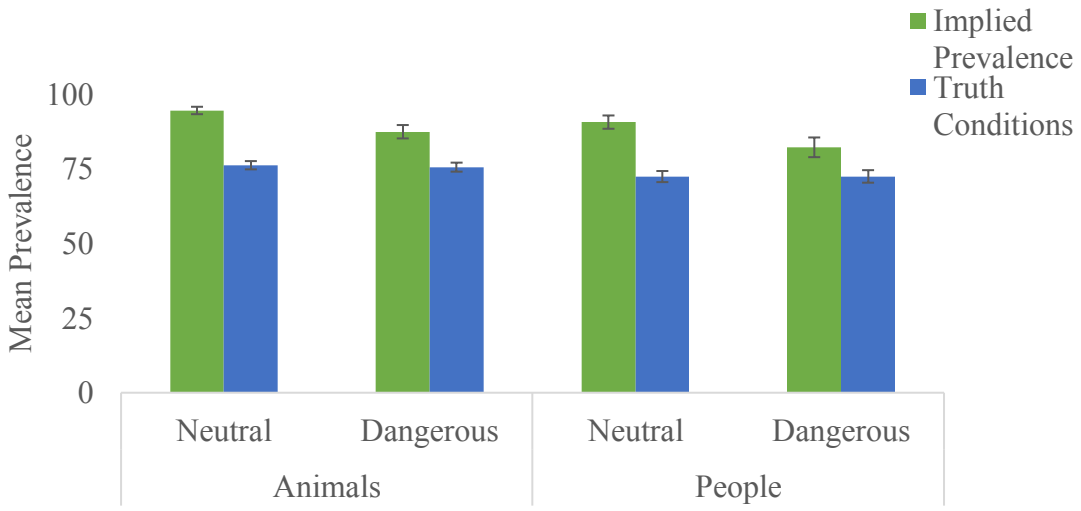


Figure 4. Study 3, the mean prevalence implied by generic statements (*Implied Prevalence* task, green bars) and the mean prevalence that led participants to accept the same generic statements (*Truth-Conditions* tasks, blue bars). Error bars represent one standard error.

Given this interaction, we examined the simple main effects of property valence within the *Implied Prevalence* and *Truth-Conditions* tasks separately. In the *Implied Prevalence* task, the mean prevalence implied by generic statements was higher for neutral properties (M = 92.7%) than dangerous properties (M = 84.9%), $F(1, 225) = 52.97$, $p < .001$, $\eta_p^2 = .19$. In contrast, in the *Truth-Conditions* task, the mean prevalence that led participants to accept the generic statements did not differ based on property valence (M_{neutral} = 74.4%; $M_{\text{dangerous}}$ = 74.1%), $F(1, 225) = 0.07$, $p = .79$, $\eta_p^2 = .00$.

We were additionally interested in whether people's responses differed depending on whether the property described a trait or a behavior. To explore the effect of property content, we examined participants' responses in the *Implied Prevalence* task (we could not conduct this analysis in the *Truth-Conditions* tasks, because each property was randomly assigned to a prevalence level, and thus participants did not necessarily receive an equal number of questions about traits/behaviors for both dangerous and neutral properties at the same prevalence levels). We conducted a repeated measures ANOVA with property valence (dangerous vs. neutral) and property content (behavior vs. trait) as within-subjects factors. We found a main effect of property valence, $F(1, 119) = 35.33$, $p < .001$, $\eta_p^2 = .23$, with higher ratings for neutral properties (M = 92.5%) than dangerous properties (M = 84.6%). We additionally found a main effect of property content, with higher ratings for traits (M = 90.0%) than behaviors (M = 87.1%), $F(1, 119) = 16.54$, $p < .001$, $\eta_p^2 = .12$. However, these main effects need to be interpreted within the context of a significant interaction, $F(1, 119) = 15.21$, $p < .001$, $\eta_p^2 = .11$. To further explore this interaction, we examined the simple main effects. When participants judged the implied prevalence of a dangerous property, their ratings were higher for traits (M = 87.5%) than behaviors (M = 81.7%), $F(1, 119) = 20.62$, $p < .001$, $\eta_p^2 = .15$. In contrast, when participants judged a neutral property, their ratings did not differ based on the content of the property (M_{traits} = 92.5%; $M_{\text{behaviors}}$ = 92.5%), $F(1,119) = 0.001$, $p = .98$, $\eta_p^2 = .000$.[9]

[9] Although the length differences in our items were minimal, we also conducted a supplementary analysis by dropping the two shortest dangerous behavior items and the two longest neutral behavior items. We replicate our results in these analyses when we control for predicate length (see "Study 3 – Length differences" in the Appendix).

As in Study 2, we examined whether there was a significant correlation between a person's score on Reynolds'(1982) social desirability scale and their mean prevalence score averaged across neutral and dangerous properties in each of the four conditions. Here we found no significant correlations.

3.4.2.3 Did domain and property valence affect generics' acceptability at different prevalence levels?

The second set of analyses examined whether responses in the *Truth-Conditions* task differed based on prevalence level. We submitted participants' responses to a logistic mixed-effects model using the glmer command in the lme4 package in R (Bates, 2007). In this model, we included domain (animals = 0; people = 1; between subject), property valence (neutral = 0; dangerous = 1; within subject), and prevalence (.1, .3, .5, .7, .9, 1; within subject) and their interactions as predictors (see Table 3). All predictors were mean centered. We also included *participant* as a random effect. We observed a main effect of prevalence, indicating that generic sentences were more likely to be judged true for higher than lower prevalence levels. However, this main effect need to be interpreted within the context of a significant domain by prevalence interaction (see Figure 5).

Fixed Effects	**Estimate**	***SE***	***p*-value**
(Intercept)	2.25	0.31	< .001
Domain	-1.12	0.60	.06
Property Valence	0.10	0.17	.55
Prevalence	10.44	0.55	< .001
Domain x Property Valence	-0.58	0.33	.08
Domain x Prevalence	-6.92	1.01	< .001
Property Valence x Prevalence	-0.21	0.59	.72
Domain x Property Valence x Prevalence	0.52	1.16	.65
Random Effect		**SD**	
Participant	Intercept	2.89	

Table 3. Logistic regression predicting "true"/ "false" judgments, based on domain, property valence, and prevalence and their interactions in Study 3.

Post-hoc tests revealed that participants were more likely to endorse social generics than generics about animals at the 10% level (Average Marginal Effect (AME) = 0.17, *SE* = 0.06, *p* = .006, 95% CI = 0.05, 0.29) (see Figure 4). At the 30% and the 50% level, there was no difference between endorsement of social generics and generics about animals (AME = 0.11, *SE* = 0.08, *p* = .16, 95% CI = -0.04, 0.26; AME = -0.07, *SE* = 0.08, *p* = .36, 95% CI = -0.22, 0.08, respectively). We observed that participants were more likely to endorse generics about animals than social generics at higher prevalence levels: 70% level (AME = -0.13, *SE* = 0.05, *p* = .01, 95% CI = -0.24, -0.03), 90% level (AME = -0.12, *SE* = 0.04, *p* = .006, 95% CI = -0.20, -0.03), and 100% level (AME = -0.09, *SE* = 0.03, *p* = .003, 95% CI = -0.15, -0.03). For a complete overview of the mean endorsement percentages at each prevalence level by domain and property valence, see "Truth-Conditions – Study 3" in the Appendix.

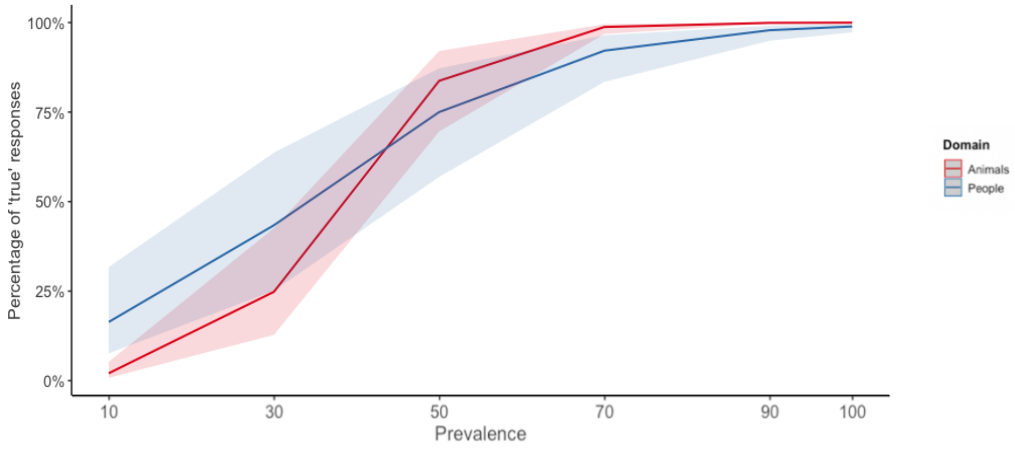


Figure 5. Mean percentage of "true" responses in Truth-Conditions task in Study 3 by domain and prevalence level.

3.4.3 Conclusions

The results of Study 3 provide additional support for the robustness and generalizability of the generic asymmetry. In particular, the inferential asymmetry elicited by social generics did not differ from that elicited by generics about animals, and this held for both neutral and dangerous properties. This suggests that the interpretation of social generics and of generics

about animals are very similar and it is consistent with previous studies that documented how generic language elicits essentialist reasoning in both domains (Foster-Hanson et al., 2016, 2019; Gelman et al., 2010; Leshin et al. 2020; Rhodes et al., 2012, 2018b).

Despite the overall lack of domain effects, we did obtain effects of valence: neutral generics yielded higher levels of implied prevalence than dangerous generics. This result differs from Cimpian et al. (2010), who reported no variation in the prevalence estimations as a function of property valence. A first possibility is that participants were reluctant to report that nearly all the members of a category present extremely cruel and harmful properties like the ones used in this study, for social desirability reasons. However, in each of the four conditions, we did not find a significant correlation between participants' score on the social desirability scale and their mean prevalence score averaged across neutral and dangerous properties.

Another possible factor could be the widespread belief that people's "true selves" are good (Strohminger, Knobe, & Newman, 2017), and perhaps this extends to novel animals as well. However, if this was the case, then we should expect to see a corresponding effect in the Truth-Conditions task. For example, participants should have found it harder to accept dangerous generic predications than neutral ones at the lower prevalence levels. On the contrary, we observed a property valence effect in the Implied Prevalence task only.

A third potential contributor to this effect may be that people have developed a set of expectations about the use of generics, based on their experience with generics concerning *familiar* dangerous versus novel properties. It has been observed that rare but dangerous properties may readily be expressed with generics (e.g., "Sharks attack swimmers"; "Mosquitoes carry the West Nile Virus"), and this may paradoxically lead to the inverse expectation that a *novel* generic about a dangerous property indicates that it is likely to be relatively rare. This possibility is also consistent with previous evidence showing that dangerous generic predications, compared to other generic predications, were interpreted as referring to common dispositions rather than prevalent properties (Prasada et al., 2013), potentially leading to the expectation that the relevant property is less generalizable

(Lazaridou-Chatzigoga et al., 2019). Thus, this seems to be the most plausible explanation for this result.

Another finding from Study 3 is that while judging the implied prevalence of dangerous generics, participants' estimates were higher for traits than behaviors. A possible explanation for this is that dangerous traits cover a wider range of threatening properties than dangerous behaviors; indeed, while there are lot of different ways to manifest a trait like "being dangerous", there are only specific ways to manifest a behavior like "hunting strangers". Thus, the former might have led to higher estimations than the latter. In contrast, when evaluating neutral generics, participants' estimates were similar for traits and behaviors. This discrepancy with the evaluation of dangerous generics might be due to a ceiling effect. As noted earlier, implied prevalence estimates for neutral properties were higher than those for dangerous properties, such that they might not have allowed observing any difference due to property content.

Although the inferential asymmetry elicited by social generics did not differ from that elicited by generics about animals, we replicated the interaction between domain and prevalence observed in the Truth-Conditions task of Study 2. Participants were more likely to accept social generics at the 10% level, whereas generics about animals were more likely to be accepted at the 70%, 90%, and 100% levels. In both Study 2 and Study 3, the interaction between prevalence level and domain was a small effect. Nevertheless, these results indicate that people were sensitive to domain at different prevalence levels in the Truth-Conditions task. As previously noted (see § 3.2.1), stereotypes operate by generalizing well beyond the evidence (e.g., Allport, 1954; Wodak et al., 2015). For this reason, social generics might be easily accepted even based on minimal evidence compared to generics about animals. Animal categories, instead, are often perceived as less variable than social categories (e.g., Brandone, 2017; Nisbett et al., 1983). As a consequence, if a small percentage of members of an animal category display a certain property, people might be less willing to accept the corresponding generalization. This possibility would also explain why, conversely, participants were less conservative in accepting generics about animals at the highest prevalence levels: since animals categories are perceived as more homogeneous than social categories, people might

be more willing to accept the corresponding generalizations about animals in case high proportions of category members share the ascribed properties.

Finally, in contrast to Cimpian et al. (2010), the participants in our study accepted dangerous and neutral generics to the same extent in the Truth-Conditions task. This discrepancy might be due to the different properties we employed. In Study 3a, we checked for this possibility by testing the same dangerous and neutral animal items used by Cimpian et al. (2010). We also investigated whether a valence effect would be observed when testing an analogous set of people items.

3.5 Study 3a

In Study 3a, we aimed at better understanding an unexpected finding of Study 3. Specifically, in contrast to Cimpian et al. (2010), we found that people rated dangerous and neutral generics alike in the Truth-Conditions task. This may have been due to the properties we used, which differed from those of Cimpian et al. (2010). In Study 3a, we examined this by testing the acceptance conditions with the same animal items as Cimpian et al. (2010), focusing exclusively on the Truth-Conditions task. We also examined whether the same pattern of responses was observed in the People condition, using an analogous set of items for groups of people. Finally, we investigated how responses in the Animal condition compared to those in the People condition. The preregistration of this study is available at https://aspredicted.org/blind.php?x=6kh579.

3.5.1 Method

3.5.1.1 Participants

One-hundred and eleven adults from the United States (48 men, 62 women, 1 non-binary/third gender; Mean Age = 41.01 years; range = 24–74 years) completed the study online and were paid $ 0.75. Participants were randomly assigned to either the *Animals–*

Truth-Conditions (n = 58) or the *People–Truth-Conditions* (n = 53) tasks. Participants were 80% White, 8% Black or African American, 6% Multiethnic, 3% Asian or Asian American, and 3% Latino or Hispanic. Seven participants were tested and excluded from the final sample because they failed the same manipulation check used in the previous studies (n = 4 in the *People–Truth-Conditions task* and n = 3 in the *Animals–Truth-Conditions*). One participant was tested and excluded from the final sample for having a non-US IP address. Finally, one participant was tested and excluded from the final sample for not being a native speaker of English.

3.5.1.2 Materials and procedure

As in Study 3, we created a list of 24 items with 12 neutral properties (e.g., "Ks have silver fur") and 12 dangerous properties (e.g., "Ks have dangerous silver fur. This fur sheds particles that get lodged in your lungs and make it impossible to breathe"; see Appendix S4 for a complete list of the items). For the animal items, we used the 10 neutral and dangerous properties from Cimpian et al. (2010), and added two additional properties for each type of property. (This was done to add two items at the 100% prevalence level, which Cimpian et al. (2010) did not include; see § 3.2). For the people items, we were concerned that ascribing the same properties to groups of people would have resulted in the perception of these groups as non-humans. For this reason, we decided to describe objects that these groups of people use, rather than describing body parts (e.g., "Ks use dangerous silver sprayers. These sprayers shed particles that get lodged in your lungs and make it impossible to breathe"). In both conditions, for each participant, the neutral and dangerous properties were randomly assigned to the 24 labels used in Experiment 2 (see Appendix S2 for a list of the labels). The order of the items was randomized for each participant.

Participants read the same introductory text as in Studies 2 and 3. All participants in this study completed the *Truth-Conditions* task. After the task, participants were also asked to complete the same manipulation check as in Studies 2 and 3. However, we did not include

Reynolds' (1982) measure of social desirability in this study, due to the lack of significant correlations observed in previous studies.

3.5.2 Results and discussion

As in the *Truth-Conditions* task of Study 3, we submitted participants' responses to a logistic mixed-effects model.[10] In this model, we included domain (animals = 0; people = 1; between subject); property valence (neutral = 0; dangerous = 1; within subject), and prevalence (.1, .3, .5, .7, .9, 1; within subject) and their interactions as predictors (see Table 4). All predictors were mean centered. We also included *participant* as a random effect. We observed a main effect for domain, indicating that participants were more likely to judge social generics than generics about animals to be true. We also observed a main effect for prevalence. For a complete overview of the mean endorsement percentages at each prevalence level by domain and property valence, see "Truth-Conditions – Study 3a" in the Appendix.

Fixed Effects	**Estimate**	***SE***	***p*-value**
(Intercept)	5.26	0.64	< .001
Domain	2.83	1.12	.01
Property Valence	0.32	0.20	.12
Prevalence	10.67	0.66	< .001
Domain x Property Valence	0.12	0.41	.77
Domain x Prevalence	0.92	1.20	.45
Property Valence x Prevalence	-0.95	0.61	.12
Domain x Property Valence x Prevalence	1.31	1.23	.29
Random Effect		**SD**	
Participant	Intercept	4.96	

Table 4. Logistic regression predicting "true"/ "false" judgments, based on domain, property valence, and prevalence and their interactions in Study 3a.

[10] In the preregistration for this study, we indicated that we would use ANOVAs to analyze the data. However, because the data are binary, it was necessary to analyze the data using non-parametric statistics; we thus opted for logistic regression.

The results of Study 3a show that participants endorsed dangerous and neutral generics at equivalent rates. Moreover, although overall the proportion of "true" responses was higher for social generics than for generics about animals, this time we did not observe an interaction between domain and prevalence level. On the one hand, these findings confirm those of Study 3, where we observed that property valence did not affect participants' responses in the Truth-Conditions task. On the other hand, these findings differ from those of Study 2 and Study 3, where we observed no domain effects and small interaction between prevalence level and domain. This discrepancy is likely due to differences in the equivalence of the items between these studies, as in Study 3a we needed to test a different set of properties across domains.

Together with the results of Study 3, these findings suggest that danger alone is not sufficient to make dangerous generic predications easier to accept; something else is required to observe this effect. Differently from our studies, Cimpian et al. (2010) compared the interpretation of neutral generics to dangerous generic predications that were either a) both dangerous and distinctive or b) intermingled with generics that varied in distinctiveness (see § 3.4). These combinations might have primed participants to think about this particular factor throughout. Thus, a possible explanation for the discrepancy between our results and Cimpian et al.'s (2010) findings is that dangerous generic predications need to be perceived as distinctive to be accepted more easily than neutral generic predications. However, further research is needed to determine whether this hypothesis is correct.

3.6 General discussion

The main goal of the present studies was to investigate the robustness of the inferential asymmetry of generics – that is, the discrepancy between people's judgments about the prevalence of a property expressed in a generic statement and the prevalence of a property among members of a category that led them to accept the sentence as being true. Specifically, we wanted to test whether the asymmetry observed for generics about animal categories extended to the interpretation of social generics. We were also interested in investigating whether the asymmetry varied as a function of property type (physical vs. non-physical) and

valence (dangerous vs. neutral). Overall, our studies confirmed that the asymmetry between generics' prevalence implications and acceptance conditions is robust: the inferential asymmetry elicited by generics about animal categories did not differ from that elicited by social generics and held for properties varying in type and valence. However, although overall property type did not affect the interpretation of generics, property valence affected generics' prevalence implications in two ways. First, neutral generic predications yielded higher prevalence estimates than dangerous generic predications. Second, when people evaluated dangerous generic predications, their prevalence estimates were higher for traits than behaviors. This discrepancy was not observed for estimates based on neutral generic predications. Finally, in the Truth-Conditions task of both Study 2 and Study 3 we observed an interaction between domain and prevalence level. Specifically, social generics were more likely to be accepted than generics about animals at the lowest prevalence levels. Generics about animals, instead, were more likely to be accepted than social generics at the highest prevalence levels.

In the present section, I discuss these results in more detail. In § 3.6.1, I consider the robustness and generalizability of the inferential asymmetry. In § 3.6.2, I review how different property types and property valences affected the interpretation of generics. In § 3.6.3, I review how domain and property valence affected generics' acceptability. In § 3.6.4, I discuss the implications of our results for the semantic of generics and for stereotyping. Finally, in § 3.6.5 I provide some concluding remarks.

3.6.1 The robustness and generalizability of the generic asymmetry

Overall, our findings provide further confirmation that the inferential asymmetry found by Cimpian et al. (2010) is robust and central to the interpretation of generics. Specifically, Study 1 showed that the generic asymmetry replicates with a larger online sample and the inclusion of the 100% level while testing generics' acceptability. Importantly, the results of Studies 2 and 3 show that the inferential asymmetry elicited by generics about animals does not differ from the one elicited by social generics.

This result was not a foregone conclusion; indeed, given that stereotypes are often accepted based on minimal evidence (e.g., Allport, 1954; Wodak et al., 2015) and that social categories are often perceived as more variable than animal categories (e.g., Brandone, 2017; Nisbett et al., 1983), we could have observed a different asymmetry for social generics, or even no asymmetry at all. On the contrary, our results suggest that the interpretation of social generics and the interpretation of generics about animals are very similar. This finding is consistent with previous evidence showing that social categories are often treated as if they were natural animal kinds (Prentice & Miller, 2007; Rhodes & Gelman, 2009; Rothbart & Taylor, 1992) and that generic language promotes essentialist reasoning in both domains (Foster-Hanson et al., 2016, 2019; Gelman et al., 2010; Leshin et al. 2020; Rhodes et al., 2012, 2018b).

As previously noted (see § 3.1.1), the generic asymmetry has been found with children as well as with adults (Brandone et al., 2015), suggesting that it is basic in human cognition. Our finding that this asymmetry extends to the interpretation of social generics provides further evidence for this hypothesis. To this end, investigating whether the generic asymmetry also extends to the interpretation of generics about artifacts and in languages other than English is an important direction for future research. Indeed, testing whether this asymmetry arises in multiple domains and languages would allow us to better understand the extent to which this phenomenon depends on fundamental cognitive processes rather than domain-specific representations and linguistic variation.

3.6.2 The effect of property type and valence on the interpretation of generics

Our results show that the generic asymmetry held for different property types (physical vs. non-physical) and property valences (dangerous vs. neutral). These results provide additional confirmation of the robustness and centrality of the generic asymmetry.

Particularly, in Study 2 we observed that the generic asymmetry did not vary as a function of property type. We hypothesize that this result might have been caused by the participants' perception of the property types used in this study as having an equivalent and high within-

category homogeneity or very similar informational values. The second possibility is consistent with Leslie's (2008) hypothesis and Cimpian et al.'s (2010) finding that, in generic form, two properties having similar informational value are evaluated alike (and differently otherwise).

In Study 3, instead, we observed that neutral generic predications yielded higher prevalence estimates than dangerous generic predications in the Implied Prevalence task. We hypothesize that this difference was due to specific expectations that people develop about the use of familiar dangerous generic predications. Generics having a dangerous valence tend to be readily accepted despite the rarity of the relevant property (e.g., "Sharks attack swimmers"); in turn, this might have led to the inverse expectation that novel dangerous generic predications have a lower implied prevalence than other novel generic predications. Consistently with this hypothesis, previous studies indicate that dangerous generic predications are expected to refer to common dispositions rather than widespread properties, leading in turn to interpret the dangerous properties as less generalizable than others (Lazaridou-Chatzigoga et al., 2019; Prasada et al., 2013).

A further finding in Study 3 is that, while evaluating dangerous generic predications, participants' estimates were higher for traits than behaviors. A possible contributor to this effect is that traits cover a wider range of properties than behaviors. For example, a trait like "being dangerous" might be manifested by numerous and different properties, while a behavior like "hunting strangers" corresponds to more specific instantiations of a single property. For this reason, traits might elicit higher estimations than behaviors. This finding joins previous evidence showing that different verbs, like "have" and "like", might have different "generalizing power" (Abelson & Kanouse, 1966; Gilson & Abelson, 1965). For example, Gilson and Abelson (1965) found that people were more likely to accept generics when they concerned category members that "have sports magazines" than when they "like sports magazines", and this might extend to the verb "are" that we included in all traits but not in behaviors. Thus, this indicates that different verbs might elicit different generic asymmetries, suggesting another possible direction for future research. However, we also observed that, while evaluating neutral generic predications, participants' estimates were not

different for traits and behaviors. We hypothesize that this discrepancy with dangerous generic predications can be traced back to a ceiling effect. Indeed, participants' estimations based on neutral properties were higher than those for dangerous properties, and this might have prevented observing any difference due to property content.

3.6.3 The effect of domain and property valence on generics' acceptability

Although the inferential asymmetry did not differ by domain, we found that people's endorsement of generics in the Truth-Conditions task differed for animal versus social categories at different prevalence levels. Specifically, in Studies 2 and 3 we observed that social generics led to a higher proportion of "true" responses than generics about animals at the lowest prevalence levels; generics about animals, instead, led to a higher proportion of "true" responses than social generics at the highest prevalence levels. We hypothesize that since stereotypes generalize beyond the available evidence (e.g., Allport, 1954; Wodak et al., 2015), social generics might be easily accepted even based on low prevalence levels. Conversely, since animal categories are often perceived as more homogeneous than social categories (e.g., Brandone, 2017; Nisbett et al., 1983), people might be less willing to accept the corresponding generalizations in case only a few members of an animal category display the same properties. Furthermore, since animal categories are often perceived as less variable, people might be more willing to accept generics about animals in case high proportions of category members share the ascribed properties.

This interaction was not observed in Study 3a; in that study, instead, we found that overall social generics were more likely to be accepted than generics about animals. We hypothesize that this discrepancy is due to differences in the equivalence of the items between these studies. While in both Study 2 and Study 3 we tested the same properties across domains, in Study 3a we needed to test different properties for people and animals.

Moreover, we did not find any tendency to endorse dangerous generics at a higher rate than neutral generics, contrary to Leslie's (2008) hypothesis and Cimpian et al.'s (2010) findings. In Study 3, we tested items that were pre-validated to test their status as dangerous

or neutral, had approximately the same number of words, and did not include generic-you expressions. In Study 3a, we tested the same neutral and dangerous animal items as Cimpian et al. (2010) together with an analogous set of people items. Our hypothesis is that danger is not sufficient per se to observe the effect theorized by Leslie (2008) and found by Cimpian et al. (2010); something extra seems to be required to increase the informational value of dangerous generic predications. Differently from our studies, Cimpian et al. (2010) compared the interpretation of neutral generics to dangerous generic predications that were either a) both dangerous and distinctive or b) intermingled with generics that varied in distinctiveness, potentially priming participants to think about *distinctiveness* throughout (see § 3.4). This discrepancy suggests that dangerous generics predications might be more likely to be accepted than neutral generic predications in case *dangerousness* is made particularly distinctive. Exploring this possibility might be important for future research, as it could better clarify the role of property valence in the interpretation of generics.

Finally, the lack of an overall interaction between property valence and domain in the Truth-Conditions task partly supports and partly contrasts with Tasimi et al.'s (2017) results. On the one hand, we replicated Tasimi et al.'s (2017) finding that threatening generic predications (e.g., "are dangerous") and non-threatening generic predications (e.g., "are helpful") are accepted alike in the social domain. Since they also did not include *distinctiveness* in their items, their result reinforces our hypothesis that this factor might play a role in leading people to accept dangerous generics more easily. On the other hand, Tasimi et al. (2017) also observed that threatening generic predications and non-threatening generic predications were not accepted alike in the artifact and animal domain. We hypothesize that this difference is due to the fact that Tasimi et al.'s (2017) non-threatening items were pro-social items, which do not have a neutral valence compared to dangerous properties. Social generics expressing pro-social properties might have different acceptance conditions than generics about animals or artifacts expressing the same properties. However, more experimental work is needed to determine whether this the case. A wider comparison between the interpretation of dangerous, distinctive, neutral, and pro-social generic predications in different domains is another important goal for future research.

3.6.4 Implications

The observation that generics about both animals and social categories are interpreted as broadly true, yet also accepted even based on low prevalence levels, provides additional support for the view that the semantics of these generalizations is not simply based on statistical prevalence (e.g., Abelson & Kanouse, 1966; Brandone et al., 2015; Carlson, 1977; Cimpian et al., 2010; Gelman et al., 2002; Krifka et al., 1995; Leslie, 2008). On the contrary, these results further reinforce the view that the interpretation of generics is based on the interaction between our conceptual knowledge and the informational value of the content being ascribed (e.g., Cimpian et al., 2010; Leslie, 2008; Prasada et al., 2013). Indeed, the results of Study 3 indicate that the prevalence estimates based on neutral generic predications were higher than those based on dangerous generic predications. Furthermore, in both Study 2 and Study 3 we observed that domain affected generics' acceptability at different prevalence levels.

Our finding that the inferential asymmetry extends to the interpretation of social generics also has important implications for stereotyping. Our results confirm that they are assumed to be broadly representative despite being accepted even based on low prevalence. In absolute terms the average prevalence level at which participants endorsed generics was high. However, considering the proportion of "true" responses at each prevalence level, social generics were often accepted based on prevalence levels lower than 50% (approximately 45% of the times; see "Truth-Conditions – Study 2", "Truth-Conditions – Study 3", and "Truth-Conditions – Study 3a" in the Appendix). When accepted at low prevalence levels, generics could contribute to the perpetuation of stereotypes. Consequently, they could be particularly pernicious and misleading in communication when used to describe social categories. Indeed, social generics are a common means to express stereotypes (Gelman et al., 2004), and the current results help explain a way in which they might contribute to the transmission and flourishing of unfounded generalizations. In previous works, several scholars assume that the observation of the generic asymmetry concerning animal categories would extend to the interpretation of social generics. Based on this assumption, they argue that social generics might profoundly affect our social cognition.

For example, both Cimpian et al. (2010) and Brandone et al. (2015) suggest that generics like "Girls are bad at math" would be easily accepted based on little evidence and despite the existence of substantial counterevidence; however, once accepted, such generics may be interpreted as referring to nearly all the category members. Consequently, they argue, generics may strongly impact on perceptions, beliefs, and behavior. Similarly, Bian and Cimpian (2017) refer to the inferential asymmetry to support their argument that stereotypes as generics beliefs about groups are not accurate descriptions of the world. Other scholars mention the results of Cimpian et al. (2010) to support the argument that the generic content of stereotypes would explain why the latter tolerate exceptions (Hammond & Cimpian, 2017) and to argue that social generics are dangerous because they are both easy to accept and hard to refute (Langton et al., 2012).

Based on Cimpian et al.'s (2010) results, I myself have discussed how generics affect our social cognition compared to quantified generalizations. In Rosola and Cella (2020), we argue that the inferential asymmetry contributes to making generics especially suited to foster the development of identity prejudice – that is, as defined by Fricker (2007), prejudice related to social identity. Fricker considers universally quantified generalizations (like "all women are intuitive") as the strongest and generalizations involving the quantifier "many" (like "many women are intuitive") as the weakest. As we point out, this is true if *predictive power* (i.e., the extent to which a generalization gives information concerning an arbitrary member of the category) is concerned, but not if *tolerance to counterevidence* is. For example, universally quantified generalizations apply to every member of the relevant category and, thus, might have pervasive effects on our perception of those individuals; however, universally quantified statements are also very easy to defeat. One single counterinstance is sufficient to falsify universally quantified generalizations, making these sentences fragile and easy to counter. The extent to which a generalization might impact on our social cognition will vary depending on both the proportion of individuals it refers to and on how hard it is to challenge and eradicate it. The inferential asymmetry shows that generics are both highly predictive (they are interpreted as applying to a high proportion of category members) and strongly tolerant to counterevidence (they are accepted as true even if few members have the

ascribed property). For this reason, we argue, they are particularly suited to give rise to pervasive and resistant stereotypes.

As discussed above, several scholars assume that the generic asymmetry extends to the interpretation of social generics. However, to date, it had not been investigated yet whether this is actually the case. Importantly, the present studies confirmed that the inferential asymmetry elicited by generics about animal categories does not differ from the one elicited by social generics. Overall, this finding increases the validity of previous work built on Cimpian et al.'s (2010) results, together with our understanding of the relation between generics and social cognition.

Finally, the present results complicate the debate concerning whether social generics are more readily accepted when ascribing dangerous properties than neutral properties. Based on Cimpian et al.'s (2010) original findings, several authors (e.g., Langton et al., 2012; Leslie, 2017) implicitly assumed that social generics ascribing dangerous properties (e.g., "Italians are mobsters") are especially pernicious because they are accepted more easily than social generics with different valences (see Chapter 2, § 2.2.3). However, while our studies confirmed the robustness and generalizability of the inferential asymmetry for social generics, it remains unclear under what circumstances dangerous generic predications might be more easily accepted than others. As discussed in the previous subsection, in future research it would be worth testing whether dangerous properties need to be perceived as distinctive to be accepted more easily than neutral properties.

3.6.5 Conclusions

These findings confirmed that generics robustly demonstrate an inferential asymmetry, in both the animal and social domains. Furthermore, our results show that the asymmetry held for different property types (physical vs. non-physical) and property valences (dangerous vs. neutral). Finally, these results show that generics' implied prevalence is affected by property valence, whereas their acceptance conditions are affected by domain at different prevalence levels. Overall, our findings further support the view that generics are a linguistic outlet of

our conceptual knowledge and that they are a powerful means to convey information about animal and social categories.

CONCLUDING REMARKS

The main focus of this dissertation was to investigate the robustness and generalizability of the inferential asymmetry for social generics. As highlighted in Chapters 1 and 2, generics have both challenging linguistic features and unique cognitive implications, which have attracted the attention of linguists, philosophers, and psychologists. In Chapter 1, I specifically focused on the main works concerning the varieties of generics, their comparison with quantified sentences and existentials, and their formal analysis. Regarding the semantics of generics, the content-based approach (e.g., Leslie, 2007, 2008) currently has the most explanatory power. While still presenting some theoretical issues, this framework accounts for our intuitions on a great variety of generic sentences. Furthermore, as noted in Chapter 2, the content-based approach is supported by several empirical confirmations and provides innovative tools to analyze the use of stereotypes expressed in generic form. The works of several scholars operating under the content-based framework illustrate how the semantic and cognitive approaches to the problems posed by generics might be combined to advance in the research fruitfully (e.g., Cimpian, Brandone, & Gelman, 2010; Prasada, Khemlani, Leslie, and Glucksberg, 2013).

In Chapter 2, instead, I focused on the cognitive implications of generics. Specifically, I first presented the hypothesis that generics are default generalizations, together with the empirical evidence motivating this hypothesis and its experimental confirmations (e.g., Leslie, 2007, 2008; Gelman, 2003, 2010; Hollander, Gelman, & Star, 2002; Leslie & Gelman, 2012). Also, I considered social generics' ability to promote essentialist beliefs (e.g., Rhodes, Leslie, & Tworek, 2012) and the implications of the content-based framework for stereotyping (e.g., Leslie, 2017), together with some controversies concerning whether social generics are especially dangerous compared to quantified generalizations (e.g., Saul, 2017). To date, there is no conclusive evidence on this matter. However, generics are a common means to describe essentialized social categories (e.g., Gelman, Taylor, & Nguyen, 2004); furthermore, they are interpreted as broadly true, yet also accepted based on weak evidence

(Cimpian et al., 2010; Brandone, Gelman, & Hedglen, 2015; the present studies). Consequently, we might have good reasons for considering social generics as more pernicious and misleading than quantified generalizations.

Finally, in Chapter 3, after reviewing several key background issues concerning the inferential asymmetry and its possible variations as a function of domain and property content, I focused on the studies I carried out with Marchak, Bianchi, and Gelman. Our findings confirm that that the asymmetry between generics' prevalence implications and acceptance conditions is robust in both the animal and social domains. Furthermore, our results show that the asymmetry holds for different property types (physical vs. non-physical) and property valences (dangerous vs. neutral). Finally, these results show that the implied prevalence of generics is affected by property valence, whereas the acceptance conditions of generics are affected by domain at different prevalence levels. Importantly, these findings lay the groundwork for future research. As discussed in Chapter 3, promising directions for further explorations include investigating whether the generic asymmetry extends to artifact categories and in languages other than English; examining whether the generic asymmetry varies as a function of the generalizing power of the verbs occurring in the generic; and further exploring under what circumstances domain and property valence affect the interpretation of generics, also widening the comparison to pro-social properties.

An additional direction for future research would be investigating the limits of the tolerance to counterevidence displayed by generics. This goal could be pursued by testing whether generics are accepted even at prevalence levels lower than 10%, or whether generics are accepted even when the existence of counterevidence is made particularly salient. For example, the prevalence level 30% in the Truth-Conditions task might be expressed using a negative framing, as in "70% of Ks *don't* have F". This methodological difference might also allow controlling whether the acceptance conditions of generics are sensitive to how the relevant prevalence levels are presented. Exploring further potential variations of the inferential asymmetry and of the acceptance conditions of generics is an important goal for future research.

Many semantic issues raised by generics are still to be solved. However, investigating the role of generics in human reasoning goes beyond the linguistic interest. Generics are ubiquitous and primary in our cognition. Our findings add to the previous literature showing that generics are a linguistic outlet of our conceptual knowledge about categories. Moreover, our results further support the view that generics are a powerful means to convey information about both animal and social categories. For this reason, these findings also help to explain a way in which generics might contribute to the transmission and flourishing of unfounded generalizations.

References

Abelson, R. P., & Kanouse, D. E. (1966). Subjective acceptance of verbal generalizations. In S. Feldman (Ed.), *Cognitive consistency: Motivational antecedents and behavioral consequents* (pp. 171–197). New York: Academic Press.

Allport, G.W. (1954). *The nature of prejudice*. Reading, MA: Addison-Wesley.

Asher, N., & Morreau, M. (1995). What some generic sentences mean. In G. Carlson & F. J. Pelletier (Eds.), *The generic book* (pp. 300–339). Chicago: Chicago University Press.

Austin, J. L. (1962). *How to do things with words*. London: Oxford University Press.

Baldwin, D.A., E. Markman, & Melartin, R. (1993). Infants' ability to draw inferences about nonobvious object properties: Evidence from exploratory play. *Child Development*, *64*(3), 711–728.

Barwise, J., & Cooper, R. (1981). Generalized quantifiers and natural language. *Linguistics and Philosophy*, *4*, 159–219.

Bates, D. (2007). *Linear mixed model implementation in lme4*. University of Wisconsin–Madison.

Bennett, M. (1996). Men's and women's self-estimates of intelligence. *The Journal of social psychology*, *136*(3), 411–412.

Bian, L., & Cimpian, A. (2017). Are stereotypes accurate? A perspective from the cognitive science of concepts. *Behavioral and Brain Sciences*, *40*, 22–24.

Bian, L., Leslie, S. J., & Cimpian, A. (2017). Gender stereotypes about intellectual ability emerge early and influence children's interests. *Science*, *355*, 389-391.

Bian, L., Leslie, S. J., Murphy, M. C., & Cimpian, A. (2018). Messages about brilliance undermine women's interest in educational and professional opportunities. *Journal of Experimental Social Psychology*, *76*, 404-420.

Bott, L., & Noveck, I. A. (2004). Some utterances are underinformative: The onset and time course of scalar inferences. *Journal of memory and language*, *51*(3), 437–457.

Brandone, A. C. (2017). Changes in beliefs about category homogeneity and variability across childhood. *Child Development*, *88*(3), 846–866.

Brandone, A. C., Cimpian, A., Leslie, S. J., & Gelman, S. A. (2012). Do lions have manes? For children, generics are about kinds rather than quantities. *Child Development*, *83*(2), 423–433.

Brandone, A. C., & Gelman S. A. (2009). Differences in preschoolers' and adults' use of generics about novel animals and artifacts: a window onto a conceptual divide. Cognition, *110*(1), 1–22.

Brandone, A. C., Gelman, S. A., & Hedglen, J. (2015). Children's developing intuitions about the truth Conditions and implications of novel generics versus quantified statements. *Cognitive Science*, *39*(4), 711–738.

Brooks, P. J., & Sekerina, I. A. (2006). Shortcuts to quantifier interpretation in children and adults. *Language Acquisition*, *13*(3), 177–206.

Burton-Roberts, N. (1976). On the generic indefinite article. *Language*, *52*, 427–448.

Burton-Roberts, N. (1977). Generic sentences and analyticity. *Studies in Language*, *1*, 155–196.

Carlson, G. N. (1977). *Reference to kinds in English*. Doctoral dissertation, University of Massachusetts, Amherst. Published in 1980 by Garland, New York.

Cartwright, R. L. (1968). Some remarks on essentialism. *The Journal of Philosophy*, *65*(20), 615-626.

Cavedon, L., & Glasbey, S. (1996). The role of context in the interpretation of generic sentences. In *Proceedings of the 10th Amsterdam Colloquium* (pp. 143–162).

Chestnut, E. K., & Markman, E. M. (2016). Are horses like zebras, or vice versa? Children's sensitivity to the asymmetries of directional comparisons. *Child Development*, *87*(2), 568-582.

Chierchia, G. (1995). Individual level predicates as inherent generics. In G. Carlson & F. J. Pelletier (Eds.), *The generic book* (pp. 176–224). Chicago: Chicago University Press.

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

Chomsky, N. (2000). *New horizons in the study of language and mind*. Cambridge: Cambridge University Press.

Christophersen, P. (1939). The Articles: A Study of Their Theory and Use in English. Copenhagen: Einar Munksgaard.

Cimpian, A. (2010). The impact of generic language about ability on children's achievement motivation. *Developmental Psychology*, *46*(5), 1333–1340.

Cimpian, A., Brandone, A. C., & Gelman, S. A. (2010). Generic statements require little evidence for acceptance but have powerful implications. *Cognitive Science*, *34*(8), 1452–1482.

Cimpian, A., & Erickson, L. C. (2012). The effect of generic statements on children's causal attributions: Questions of mechanism. *Developmental Psychology*, *48*(1), 159–170.

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

Cimpian, A., & Markman, E. M. (2008). Preschool children's use of cues to generic meaning. *Cognition*, *107*(1), 19–53.

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

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

Cimpian, A., & Salomon, E. (2014). The inherence heuristic: An intuitive means of making sense of the world, and a potential precursor to psychological essentialism. *Behavioral and Brain Sciences*, *37*, 461–480.

Cohen, A. (1996). *Think generic! The meaning and use of generic sentences*. Doctoral dissertation, Carnegie Mellon University.

Cohen, A. (1999). Generics, frequency adverbs and probability. *Linguistics and Philosophy*, *22*(3), 221–253.

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

Cohen, A. (2004). Generics and Mental Representation. *Linguistics and Philosophy*, *27*(5), 529–556.

Cohen, A., & Erteschik-Shir, N. (2002). Topic, focus, and the interpretation of bare plurals. Natural Language Semantics, *10*(2), 125–165.

Coleman, J. M., & Hong, Y. Y. (2008). Beyond nature and nurture: The influence of lay gender theories on self-stereotyping. *Self and Identity*, *7*(1), 34-53.

Conrad, B. (1982). *Referring and non-referring phrases: A study in the use of the gerund and the infinitive*. Copenhagen: Akademisk Forlag.

Dahl, Ö. (1975). On generics. In E. L. Keenan (Ed.), *Formal semantics of natural language: Papers from a colloquium sponsored by the King's College Research Centre, Cambridge* (pp. 99–111). Cambridge, England: Cambridge University Press.

Declerck, R. (1986). The manifold interpretations of generic sentences. *Lingua*, *68*, 149–188.

Declerck, R. (1988). Restrictive "when"-clauses. *Linguistics and Philosophy*, *11*, 131-168.

Declerck, R. (1991). The origins of genericity. *Linguistics*, *29*(1), 79–102.

Everett, D. L. (2005). Cultural constraints on grammar and cognition in Pirahã. *Current Anthropology*, 46, 621–646.

Everett, D. L. (2009). Pirahã culture and grammar: A response to some criticisms. *Language*, 85(2), 405-442.

Farkas D. (1981), Intensionality and Romance Subjunctive Relatives. Doctoral dissertation distributed by Indiana University Linguistics Club, Bloomington, Indiana.

Farkas, D., & Sugioka, Y. (1983), Restrictive if/when clauses. *Linguistics and Philosophy*, 6, 225–258.

Fernald, T. (2000). *Predicates and temporal arguments*. Oxford: Oxford University Press.

Foster-Hanson, E., Leslie, S. J., & Rhodes, M. (2016). How does generic language elicit essentialist beliefs? In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 1541-1546). Cognitive Science Society.

Foster-Hanson, E., Leslie, S. J., & Rhodes, M. (2019). Speaking of kinds: How generic language shapes the development of category representations. *PsyArXiv*. Doi:10.31234/osf.io/28qf7

Frable, D. E. S., & Bem, S. L. (1985). If you are gender schematic, all members of the opposite sex look alike. *Journal of Personality and Social Psychology*, *49*, 459–468.

Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. New York: Oxford University Press.

Gaunt, R. (2006). Biological essentialism, gender ideologies, and role attitudes: What determines parents' involvement in child care. *Sex Roles*, *55*, 523–533.

Gelman, S. A. (1988). The development of induction within natural kind and artifact categories. *Cognitive Psychology*, *20*, 65–95.

Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. Oxford, England: Oxford University Press.

Gelman, S. A. (2010). Generics as a window onto young children's concepts. In F. J. Pelletier (Ed.), *Kinds, things, and stuff: The cognitive side of generics and mass terms*. New directions in cognitive science (pp. 100–120) (Vol. 12). New York: Oxford University Press.

Gelman, S. A., & Bloom, P. (2007). Developmental changes in the understanding of generics. *Cognition*, *105*(1), 166–183.

Gelman, S. A., Cimpian, A., & Roberts, S. O. (2018). How deep do we dig? Formal explanations as placeholders for inherent explanations. *Cognitive Psychology*, *106*, 43–59.

Gelman, S. A., Coley, J. D., Rosengren, K. S., Hartman, E., Pappas, A, & Keil, F.C. (1998). Beyond labeling: The role of parental input in the acquisition of richly-structured categories. *Monographs of the Society for Research in Child Development*, *63*, Serial No. 253.

Gelman, S. A., Goetz, P. J., Sarnecka, B. W., & Flukes, J. (2008). Generic language in parent-child conversations. *Language Learning and Development*, *4*, 1–31.

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

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

Gelman, S. A., & Markman, E. M. (1986). Categories and induction in young children. *Cognition*, *23*, 183–209.

Gelman, S. A., & Markman, E. M. (1987). Young children's inductions from natural kinds: The role of categories and appearances. *Child Development*, *58*, 1532–1541.

Gelman, S. A., & Raman, L. (2003). Preschool children use linguistic form class and pragmatic cues to interpret generics. *Child Development*, *74*, 308–325.

Gelman, S. A., Star, J. R., & Flukes, J. E. (2002). Children's use of generics in inductive inferences. *Journal of Cognition and Development*, *3*, 179–199.

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

Gelman, S. A., Taylor, M. G., & Nguyen, S. P. (2004). Mother-child conversations about gender: Understanding the acquisition of essentialist beliefs. *Monographs of the Society for Research in Child Development*, *69*, Serial No. 275.

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

Gelman, S. A., & Wellman, H. M. (1991). Insides and essences: Early understandings of the non-obvious. *Cognition*, *38*, 213–244.

Geurts, B. (1985). Generics. *Journal of Semantics*, *4*(3), 247–255.

Gilson, C., & Abelson, R. P. (1965). The subjective use of inductive evidence. *Journal of Personality and Social Psychology*, *2*(3), 301–310.

Goldin-Meadow, S. (2003). *The resilience of language*. New York: Psychology Press.

Goldin-Meadow, S., Gelman, S. A., & Mylander, C. (2005). Expressing generic concepts with and without a language model. *Cognition*, *96*(2), 109–126.

Graham, S. A., Kilbreath, C. S., & Welder, A. N. (2001). Words and shape similarity guide 13–month-olds' inferences about nonobvious object properties. In J. D. Moore, & K. Stenning (Eds.), *Proceedings of the twenty third annual conference of the Cognitive Science Society* (pp. 352–357). Hillsdale, N.J.: Erlbaum.

Graham, S. A., Nayer, S. L., & Gelman, S. A. (2011). Two-year-olds use the generic/nongeneric distinction to guide their inferences about novel kinds. *Child Development*, *82*(2), 493-507.

Greenberg, Y. (2003). *Manifestations of genericity*. New York: Routledge.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics*, Vol. 3 (pp. 41–58). New York: Academic Press.

Halberda, J., Taing, L., & Lidz, J. (2008). The development of 'most' comprehension and its potential dependence on counting ability in preschoolers. *Language Learning and Development*, *4*, 99–121.

Hammond, M. D., & Cimpian, A. (2017). Investigating the cognitive structure of stereotypes: Generic beliefs about groups predict social judgments better than statistical beliefs. *Journal of Experimental Psychology: General*, *146*(5), 607-614.

Haslam, N., Rothschild, L., & Ernst, D. (2000). Essentialist Beliefs About Social Categories. *British Journal of Social Psychology*, *39*: 113–127.

Haslanger, S. (2011). Ideology, generics, and common ground. In C. Witt (Ed.), *Feminist metaphysics: Explorations in the ontology of sex, gender, and the self* (pp. 179–209). Dordrecht: Springer.

Haslanger, S. (2014). The normal, the natural and the good: Generics and ideology. *Politica & Società*, *3*(3), 365–392.

Heim, I. (1982). *The semantics of definite and indefinite noun phrases*. Doctoral dissertation, University of Massachusetts, Amherst.

Heyer, G. (198). Generic descriptions, default reasoning, and typicality. *Theoretical Linguistics*, *12*(1), 33–72.

Hoicka, E., Saul, J. M., Prouten, E., Whitehead, L., & Sterken, R. (ms). Language signalling high proportions, not just generics, leads to essentializing for novel social kinds. *PsyArXiv*. Doi:10.31234/osf.io/xe6sj

Hollander, M. A., Gelman, S. A., & Star, J. (2002). Children's interpretations of generic noun phrases. *Developmental Psychology*, *38*(6), 883–894.

Jespersen, O. (1927). *A modern English grammar*. Heidelberg: Carl Winter's Universitätsbuchhandlung.

Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In: Gilovich, T., Griffin, D., & Kahneman, D. (Eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment* (pp. 49–81). New York: Cambridge University Press.

Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). Oxford: Oxford University Press.

Karczewski, D., Wajda, E., & Poniat, R. (2020). Do all storks fly to Africa? Universal statements and the generic overgeneralization effect. *Lingua*, *246*, 1–15.

Kay, A. C., Day, M. V., Zanna, M. P., & Nussbaum, A. D. (2013). The insidious (and ironic) effects of positive stereotypes. *Journal of Experimental Social Psychology*, *49*(2), 287-291.

Khemlani, S., Leslie, S. J., Glucksberg, S., & Rubio-Ferna´ndez, P. R. (2007). Do ducks lay eggs? How people interpret generic assertions. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th annual Cognitive Science Society* (pp. 395–400). Austin, TX: Cognitive Science Society.

Kratzer, A. (1995). Stage-level and individual-level Predicates. In G. Carlson & F. J. Pelletier (Eds.), *The generic book* (pp. 125–176). Chicago: Chicago University Press.

Krifka, M., partly in collaboration with Gerstner C. (1987). *An Outline of Genericity*. Technical Report SNS-Bericht 87-25, Seminar für natürliche Systeme, Tübingen University.

Krifka, M. (2013). Definitional generics. In A. Mari, C. Beyssade, & F. Del Prete (Eds.), *Genericity* (pp. 372–389). Oxford: Oxford University Press.

Krifka, M., Pelletier, F. J., Carlson, G. N., ter Meulen, A., Chierchia, G., & Link, G. (1995). Genericity: An introduction. In G. Carlson & F. J. Pelletier (Eds.), *The generic book* (pp. 1–124). Chicago: Chicago University Press.

Laca, B. (1990). Generics objects: Some more pieces of the puzzle. *Lingua*, *81*, 25–46.

Langton, R., Haslanger, S., & Anderson, L. (2012). Language and Race. In G. Russell & D. Graff Fara (Eds.), *The Routledge Companion to the Philosophy of Language* (pp. 753–767). London: Routledge.

Lawler, J. M. (1973). Studies in English generics. *University of Michigan Papers in Linguistics*, *1*(1), 1–184. Ann Arbor: University of Michigan Press.

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

Lazaridou-Chatzigoga, D., Katsos, N., & Stockall, L. (2017). Generic and universal generalisations: Contextualising the 'Generic Overgeneralisation' effect. In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society* (pp. 724–729).

Lazaridou-Chatzigoga, D., Katsos, N., & Stockall, L. (2019). Generalising about striking properties: do glippets love to play with fire? *Frontiers in Psychology*, *10*,1971, 1–14.

Lazaridou-Chatzigoga, D., & Stockall, L. (2013). Genericity, exceptions and domain restriction: Experimental evidence from comparison with universals. In E. Chemla, V. Homer, & G. Winterstein (Eds.), *Proceedings of Sinn und Bedeutung 17* (pp. 325–343). École Normale Supérieure, Paris.

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

Leshin, R., Leslie, S. J., & Rhodes, M. (2020). Does it matter how we speak about social kinds? A large, pre-Registered, online experimental study of how language shapes the development of essentialist beliefs.

PsyArXiv. Doi:10.31234/osf.io/nb6ys

Leslie, S. J. (2007). Generics and the structure of the mind. *Philosophical Perspectives*, *21*(1), 375–405.

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

Leslie, S. J. (2012a). Generics. In G. Russell & D. Graff Fara (Eds.), *The Routledge Companion to the Philosophy of Language* (pp. 355–367). London: Routledge.

Leslie, S. J. (2012b). Generics articulate default generalization. *Recherches Linguistiques de Vincennes: New Perspectives on Genericity at the Interfaces*, *41*, 25–45.

Leslie, S. J. (2014). Carving up the social world with generics. In T. Lombrozo, J. Knobe & S. Nichols (Eds.), *Oxford studies in experimental philosophy*, (pp. 208–232). Oxford: Oxford University Press.

Leslie, S. J. (2015). Hillary Clinton is the only man in the Obama administration: Dual character concepts, generics, and gender. *Analytic Philosophy*, *56*(2), 111–141.

Leslie, S. J. (2017). The original sin of cognition: Fear, prejudice, and generalization. *The Journal of Philosophy*, 114(8), 393-421.

Leslie, S. J., Cimpian, A., Meyer, M., & Freeland, E. (2015). Expectations of brilliance underlie gender distributions across academic disciplines. *Science*, *347*, 262-265.

Leslie, S. J., & Gelman, S. A. (2012). Quantified statements are recalled as generics: Evidence from preschool children and adults. *Cognitive Psychology*, *64*(3), 186–214.

Leslie, S. J., Khemlani, S., & Glucksberg, S. (2011). Do all ducks lay eggs? The generic overgeneralization effect. *Journal of Memory and Language*, *65*(1), 15–31.

Leslie, S. J., & Lerner, A. (2016). Generic generalizations. In E. D. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2016 edition).

Levinson, S. C. (1983). *Pragmatics*. Cambridge, England: Cambridge University Press.

Lewis, D. (1975). Adverbs of quantification. In E. Keenan (Ed.), *Formal semantics of natural language* (pp. 3–15). Cambridge, England: Cambridge University Press.

Liebesman, D. (2011). Simple generics. *Noûs*, *45*(3), 409–442.

Mandalaywala, T. M., Amodio, D. M., & Rhodes, M. (2018). Essentialism promotes racial prejudice by increasing endorsement of social hierarchies. *Social Psychological and Personality Science*, *9*(4), 461-469.

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

Mari, A., Beyssade, C., & Del Prete, F. (2013). *Introduction*. In A. Mari, C. Beyssade, & F. Del Prete (Eds.), *Genericity* (pp. 1–92). Oxford: Oxford University Press.

Martin, C. L., & Parker, S. (1995). Folk theories about sex and race differences. *Personality and Social Psychology Bulletin*, *21*(1), 45-57.

Maurer, P., Meeuwis, M.; APiCS Consortium. (2013). Generic noun phrases in subject function. In S. M. Michaelis, P. Maurer, M. Haspelmath, & M. Huber (Eds.), The Atlas of Pidgin and Creole Language Structures (pp. 114–117). Oxford: Oxford University Press.

McNeil, B. J., Pauker, S. G., Sox, H. C., Jr., & Tversky, A. (1982). On the elicitation of preferences for alternative therapies. *New England Journal of Medicine*, *306*(21), 1259-1262.

Medin, D. L., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 179-195). Cambridge: Cambridge University Press.

Meyer, M., Gelman, S. A., & Stilwell, S. M. (2011). Generics are a cognitive default: Evidence from sentence processing. In L. Carlson, C. Holscher, & T. Shipley (Eds.), *Proceedings of the 33rd annual conference of the cognitive science society* (pp. 913–918). Boston, MA: Cognitive Science Society.

Milsark, G. (1974). *Existential Sentences in English*. Doctoral dissertation, MIT.

Nickel, B. (2008). Generics and the Way of Normality. *Linguistics and Philosophy*, *31*(6), 629–648.

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

Nisbett, R. E., Krantz, D. H., Jepson, C., & Kunda, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review*, *90*(4), 339–363.

Nunberg, G., & Pan, C. (1975). Inferring Quantification in Generic Sentences. In: *Papers from the Eleventh Regional Meeting of the Chicago Linguistic Society* (pp. 412–422), Chicago: Chicago Linguistic Society.

Orvell, A., Kross, E., & Gelman, S. A. (2017). How "you" makes meaning. *Science*, *355*, 1299–1302.

Orvell, A., Kross, E., & Gelman, S. A. (2019). "You" and "I" in a foreign land: The persuasive force of generic-you. *Journal of Experimental Social Psychology*, *85*, 1–9.

Orvell, A., Kross, E., & Gelman, S. A. (2020). "You" speaks to me: Effects of generic-you in creating resonance between people and ideas. *Proceedings of the National Academy of Sciences, 117*(49), 31038-31045.

Papafragou, A., & Schwarz, N. (2006). Most wanted. *Language Acquisition: A Journal of Developmental Linguistics* [Special issue: On the Acquisition of Quantification], *13*, 207–251.

Pelletier, F. J., & Asher, N. (1997). Generics and defaults. In J. van Benthem & A. ter Meulen (Eds.), *Handbook of logic and language* (pp. 1125–1179). Cambridge, MA: MIT Press.

Pettigrew, T. F., & Tropp, L. R. (2000). Does intergroup contact reduce prejudice: Recent meta-analytic findings. In S. Oskamp (Ed.), *Reducing prejudice and discrimination* (pp. 93–114). Mahwah, NJ: Erlbaum.

Platteau, F. (1980). Definite and indefinite generics. In J. van der Auwera (Ed.), *The semantics of determiners* (pp. 112–123) London: Croom Helm.

Prasada, S. (2000). Acquiring generic knowledge. *Trends in Cognitive Sciences*, *4*(2), 66–72.

Prasada, S., & Dillingham, E. M. (2006). Principled and statistical connections in common sense conception. *Cognition*, *99*(1), 73–112.

Prasada, S., & Dillingham, E. M. (2009). Representation of principled connections: A window onto the formal aspect of common sense conception. *Cognitive Science*, *33*(3), 401–448.

Prasada, S., Khemlani, S., Leslie, S. J., & Glucksberg, S. (2013). Conceptual Distinctions Amongst Generics. *Cognition*, *126*(3), 405–422.

Prentice, D. A., & Miller, D. T. (2007). Psychological essentialism of human categories. *Current Directions in Psychological Science*, *16*(4), 202–206.

Radden, G. (2009). Generic reference in English: A metonymic and conceptual blending analysis. In K. U. Panther, L. L. Thornburg, & A. Barcelona (Eds.), *Metonymy and metaphor in grammar* (Vol. 25) (pp. 199-228). John Benjamins Publishing.

Reynolds, W.M. (1982). Development of reliable and valid short forms of the Marlowe-Crowne Social Desirability Scale. *Journal of Clinical Psychology*, *38*, 119-125.

Rhodes, M., & Gelman, S. A. (2009). A developmental examination of the conceptual structure of animal, artifact, and human social categories across two cultural contexts. *Cognitive Psychology*, *59*(3), 244–274.

Rhodes, M., Gelman, S. A., & Karuza, J. C. (2014). Preschool ontology: The role of beliefs about category boundaries in early categorization. *Journal of Cognition and Development*, *15*(1), 78-93.

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

Rhodes, M., Leslie, S. J., Saunders, K., Dunham, Y., & Cimpian, A. (2018b). How does social essentialism affect the development of inter-group relations? *Developmental Science*, *21*, 1–15.

Rhodes, M., Leslie, S. J., & Tworek, C. M. (2012). Cultural transmission of social essentialism. *Proceedings of the National Academy of Sciences*, *109*(34), 13526–13531.

Rhodes, M., & Mandalaywala, T. M. (2017). The development and developmental consequences of social essentialism. *Wiley Interdisciplinary Reviews: Cognitive Science*, *8*(4), 1–18.

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

Roberts, S. O., Ho, A. K., Rhodes, M., & Gelman, S. A. (2017). Making boundaries great again: Essentialism and support for boundary-enhancing initiatives. *Personality and Social Psychology Bulletin*, *43*(12), 1643-1658.

Rosch, E.H. (1978). Principles of categorization. In E. H. Rosch & B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, N.J.: Erlbaum Associates.

Rosch, E. H., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*(4), 573-605.

Rosola, M., & Cella, F. (2020). Generics and epistemic injustice. *Ethical Theory and Moral Practice*, *23*(5), 739–754.

Rothbart, M., & Taylor, M. (1992). Category labels and social reality: Do we view social categories as natural kinds? In G. Semin & K. Fiedler (Eds.), *Language, interaction and social cognition* (pp. 11–36). London: Sage.

Salomon, E., & Cimpian, A. (2014). The inherence heuristic as a source of essentialist thought. *Personality and Social Psychology Bulletin*, *40*(10), 1297-1315.

Sanabria, T., & Penner, A. (2017). Weeded out? Gendered responses to failing calculus. *Social Sciences*, *6*(2), 1–14.

Saul, J. M. (2017). Are generics especially pernicious? *Inquiry*, 1–18.

Schubert L. & F. Pelletier (1989), Generically speaking or using discourse representation theory to interpret generics. In G. Chierchia, B. Partee, & R. Turner (Eds.), *Properties, types and meaning* (Vol.2) (pp. 193-268). Dordrecht: Kluwer.

Segall, G., Birnbaum, D., Deeb, I., & Diesendruck, G. (2015). The intergenerational transmission of ethnic essentialism: How parents talk counts the most. *Developmental Science*, *18*(4), 543-555.

Semmerling, T. J. (2006). "Evil" Arabs in American Popular Film: orientalist fear. Austin, T.X.: University of Texas Press.

Shaheen, J. G. (2003). Reel bad Arabs: How Hollywood vilifies a people. *The ANNALS of the American Academy of Political and Social Science*, *588*(1), 171-193.

Stalnaker, R. C. (1970). Pragmatics. *Synthese*, 22, 272–289.

Stanley, J., & Szabó, Z. G. (2000). On quantifier domain restriction. *Mind and Language*, *15*, 219–261.

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and brain sciences*, *23*(5), 645-665.

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

Strohminger, N., Knobe, J., & Newman, G. (2017). The true self: A psychological concept distinct from the self. *Perspectives on Psychological Science*, *12*(4), 551-560.

Stump, G. (1985). *The Semantic Variability of Absolute Constructions*. Dordrecht: Reidel.

Sullivan, J., Bale, A., & Barner, D. (2018). Most preschoolers don't know most. *Language Learning and Development*, *14*(4), 320-338.

Sutherland, S. L., & Cimpian, A. (2015). Children show heightened knew-it-all-along errors when learning new facts about kinds: Evidence for the power of kind representations in children's thinking. *Developmental Psychology*, *51*(8), 1115-1130.

Sutherland, S. L., Cimpian, A., Leslie, S. J., & Gelman, S. A. (2015). Memory errors reveal a bias to spontaneously generalize to categories. *Cognitive science*, *39*(5), 1021-1046.

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

Tasimi, A., Gelman, S. A., Cimpian, A., & Knobe, J. (2017). Differences in the evaluation of generic statements about human and non-human categories. *Cognitive Science*, *41*(7), 1934-1957.

Tessler, M., & Goodman, N. (2019). The language of generalization. *Psychological Review*, *126*(3), 395–436.

Vasilyeva, N., & Ayala López, S. (2019). Structural thinking and epistemic injustice. In B. R. Sherman, & S. Goguen (Eds.), *Overcoming epistemic injustice: social and psychological perspectives* (Ser. Collective studies in knowledge and society) (pp. 63 - 85). Rowman & Littlefield International.

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

Vasilyeva, N., & Lombrozo, T. (2020). Structural thinking about social categories: Evidence from formal explanations, generics, and generalization. *Cognition*, 204, 1–14.

Vendler, Z. 1967. *Linguistics in Philosophy*. Ithaca, N.Y.: Cornell University Press.

Waxman, S. R. (2010). Names will never hurt me? Naming and the development of racial and gender categories in preschool-aged children. *European Journal of Social Psychology*, *40*(4), 593-610.

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

Webography

Arva, B., Idris, M., & Pervez, F. (2017). Almost all news coverage of the Barcelona attack mentioned terrorism. Very little coverage of Charlottesville did. *The Washington Post*. Available at: https://www.washingtonpost.com/news/monkey-cage/wp/2017/08/31/almost-all-news-coverage-of-the-barcelona-attack-mentioned-terrorism-very-little-coverage-of-charlottesville-did/

Butler, A. (2015). Shooters of color are called 'terrorists' and 'thugs.' Why are white shooters called 'mentally ill'? *The Washington Post*. Available at: https://www.washingtonpost.com/posteverything/wp/2015/06/18/call-the-charleston-church-shooting-what-it-is-terrorism/

Cox, A. (2004). U.S. poised for epidemic West Nile year. *CNN*. Available at: http://edition.cnn.com/2004/HEALTH/05/03/wnv.outlook/

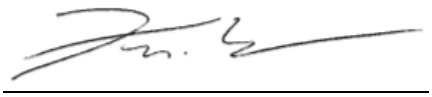
Dagli, C. K. (2015). Should Muslims really apologize for terror attacks? *CNN*. Available at: https://edition.cnn.com/2015/02/06/opinion/muslims-apology-dagli/index.html

International Shark Attack File (2014). Florida Museum of natural History, University of Florida. Available at: https://www.floridamuseum.ufl.edu/shark-attacks/odds/compare-risk/dog-attack/

Jones, S.G., Doxsee, C., & Harrington, N. (2020). The escalating terrorism problem in the United States. *Center for Strategic and International Studies*. Available at: https://www.csis.org/analysis/escalating-terrorism-problem-united-states/

Scheviz, T. (2002). FBI sees leap in anti-Muslim hate crimes / 9/11 attacks blamed for bias - blacks still most frequent victims. *San Francisco Chronicle*. Available at: http://www.sfgate.com/news/article/FBI-sees-leap-in-anti-Muslim-hate-crimes-9-11-2750152.php

All URLs last accessed on December 30th, 2020



Appendix

Study 1

In addition to the Truth-Conditions and the Implied Prevalence tasks, we also conducted a direct replication of the Truth-Conditions task from Cimpian et al. (2010). In the *Expanded Truth-Conditions* (ETC) task reported in Chapter 3 (see § 3.2), we included two items at the 100% prevalence level; however, this differed from the *Original Truth-Conditions* (OTC) task in Cimpian et al.'s (2010) studies. The results from the OTC without a 100% prevalence level are reported below.

Method

Participants. Fifty-nine adults from the United States (32 men, 27 women; Mean age = 39.46 years; range = 24–71 years) completed the study online via MTurk and were paid 40 cents. Participants were 76% White, 10% Asian or Asian American, 5% Black or African American, 3% Middle Eastern or North African, 2% American Indian or Alaska Native 2% Latino or Hispanic, and 2% Multiethnic. One additional participant was tested and excluded from the final sample for having a non-US IP address. Two other participants were tested and excluded for having duplicate IP addresses.

Materials and procedure. In the OTC task, participants judged whether generic statements were "true" based on the prevalence level of an ascribed property – each of the following prevalence levels was presented twice: 10%, 30%, 50%, 70%, 90%. Participants in the OTC task received 10 items randomly selected from the list of 12 items (see Appendix S1 for a complete list of the items).

Results and discussion

Data coding. We coded data in the OTC task in a similar way to the ETC task reported in Chapter 3. However, in the OTC task, we imputed a score of 100% to participants who said "false" to all questions.

We observed an asymmetry between ratings in the *Implied Prevalence* task and the *Original Truth-Conditions* task, with participants providing higher mean ratings in the *Implied Prevalence* task (M = 95.6%) than in the OTC task (M = 67.7%), $t(112) = 9.22$, $p < .001$, $d = 1.73$.

To examine whether participants' responses differed between the ETC and OTC tasks, we explored participants' "true/false" responses. We submitted responses to a logistic mixed-effects model using the glmer command in the lme4 package in R (Bates, 2007).[1] In this model, we included task (OTC = 0; ETC = 1; between subject), prevalence (.1, .3, .5, .7, .9; within subject), and their interaction as predictors (see Table S1). All predictors were mean centered. We also included *participant* as a random effect.[2]

Fixed Effects	**Estimate**	***SE***	***p*-value**
(Intercept)	7.65	1.83	< .001
Task	6.35	2.09	.002
Prevalence	12.24	1.24	< .001
Task x Prevalence	1.28	2.12	.55
Random Effect		**SD**	
Participant	Intercept	9.99	

Table S1. Logistic regression predicting "true" judgments, based on version of Truth-Conditions task and prevalence of the property in Study 1.

We observed a main effect of task, with ratings higher in the ETC than in the OTC task. We additionally observed a main effect of prevalence. However, there was no interaction between these factors, suggesting that participants in both versions of the task responded to items at a given prevalence level in similar ways. We thus decided to include a 100% level in subsequent tasks, because it was easier to interpret the data if we did not have to impute scores to participants who said "no" to all questions. For a complete overview of the mean

[1] In the preregistration for this study, we indicated that we would use ANOVA to analyze the data. However, since the data is binary, it is necessary to analyze the data using non-parametric statistics, such as logistic regression.

[2] We additionally fit a model including item as a random effect; however, we found that the estimate for the SD of the intercept for item was zero, so we omitted item as a random effect in the final model.

endorsement percentages at the 10-90% levels in the ETC and OTC tasks, see "Truth-Conditions – Study 1" in this appendix.

Study 2

In Study 2, we did not observe an effect of property type or domain. To observe whether the lack of effect was due to the properties being judged as having high within-category homogeneity, we conducted a follow-up study. We adapted a procedure used in previous research to observe the prevalence in a category that participants judged our properties to have (Nisbett, Krantz, Jepson, & Kunda, 1983).

Method

Participants. Seventy-four adults from the United States (28 men, 46 women; M = 40.43 years; range = 27–67 years) completed the study online and were paid 75 cents. Participants were 78% White, 11% Black or African American, 5% Multiethnic, 4% Asian or Asian American, and 1% Latino or Hispanic. Participants were randomly assigned to either the *Animal Condition* (N = 40) or the *People Condition* (N = 34). Five participants were tested and excluded from the final sample because they failed the same manipulation check used in Study 2 (4 in the *People Condition*, 1 in the *Animal Condition*). One participant was also tested and excluded from the final sample for not being a native speaker of English. We randomly presented the same attention checks items used in the previous studies, but no participant failed both of them.

Materials and procedure. Participants completed a *Prevalence Estimation* task, where they estimated the prevalence of the properties based on three members of a novel category. The structure of the items was adapted from Nisbett et al. (1983). The following is the introductory text-line we used for this study:

> "Imagine that you are an explorer who has landed on a remote island. This island is very large, and has many different **[animals]**/**[people]** on it. It is roughly the size

of Alaska, and has a lot of geographical, climatic, and environmental variety. You encounter a number of new **[animals]**/**[people]**. You observe the properties of your "samples" and you need to make guesses about how common these properties would be in other **[animals]**/**[people]**."

Each participant evaluated the same 24 properties used in Study 2, randomly paired with the same novel category labels. The order of the items was randomized for each participant. See Table S2 for example of an item as it appeared to participants.

Information: Suppose you encounter three STADES. These three STADES bury their leftover food. Question: What percentage of STADES on the island do you expect to bury their leftover food?

Table S2. Sample item from follow-up of Study 2.

Results and discussion

We conducted a repeated measures ANOVA with domain (animals vs. people) as between-subjects factor, and property type (physical vs. non-physical) as a within-subjects factor (see Figure S1). The ANOVA revealed a main effect of domain, with ratings overall higher in the *Animal Condition* (M = 86.5%) than the *People Condition* (M = 75.8%), $F(1,72) = 6.24$, p = .015, $\eta p^2 = .08$. In contrast, there was a non-significant main effect of property type ($M_{\text{non-physical}}$ = 82.4%; M_{physical} = 79.9%), $F(1,72) = 3.87$, p = .053, $\eta p^2 = .05$ and a non-significant interaction between property type and domain ($M_{\text{animal non-physical}}$ = 88.6%; $M_{\text{people non-physical}}$ = 76.2%) ($M_{\text{animal physical}}$ = 84.5%; $M_{\text{people physical}}$ = 75.3%), $F(1,72) = 1.60$, p = .21, $\eta p^2 = .02$. These data suggest that the lack of a property effect in Study 2 may be due to the properties having equivalent and high within-category homogeneity. However, it is not clear from these

results why there was no domain effect in the primary study, since we observed that properties in our studies were judged more homogenous in animal than people categories.

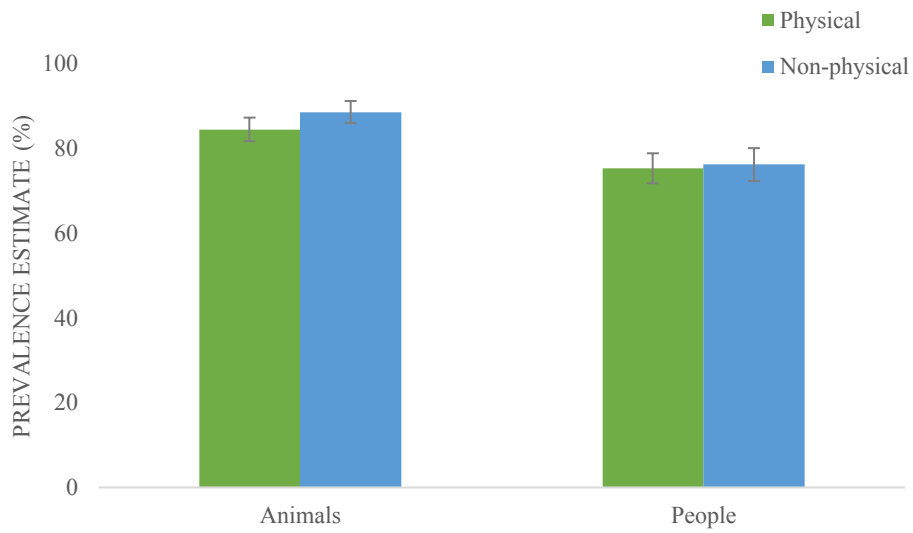


Figure S1. Study 2 follow-up, the mean prevalence estimates. Error bars represent 1 SE.

Study 3

We found that recruitment on MTurk slowed considerably after Study 2. To allow us to continue with data collection, we began recruiting non-Masters approximately half way through data collection. However, when we examined the data, we observed differences between the pattern of responses across samples (Masters vs. non-Masters). To ensure that our recruitment method was consistent across studies, we replaced the non-Masters in our sample in Study 3 with Masters. Below we report the results of our primary analyses for the non-Masters sample.

Method

Participants. One-hundred and one adults from the United States (62 men, 39 women; Mean age = 35.11 years; range = 21–64 years) completed the study online and were paid $1 for

participating. Participants were randomly assigned to one of four conditions: *Animals–Truth-Conditions* (n = 35), *People–Truth-Conditions* (n = 18), the *Animals–Implied Prevalence* (n = 24), and *People–Implied Prevalence* (n = 24). Unlike the participants reported in Chapter 3, these participants had not been granted a Master Worker qualification from MTurk. Participants were 79% White; 9% Black or African American; 5% Multiethnic; 4% Latino or Hispanic; 2% Asian or Asian American, and 1% not listed. Fifteen participants were tested and excluded from the final sample because they failed the manipulation check (n = 6 in the *People–Truth-Conditions,* n = 9 in the *People–Implied Prevalence*). Two participants were tested and excluded from the final sample for having non-US IP addresses. Finally, two participants were tested and excluded from the final sample because they were not native speakers of English.

Materials and procedure. The materials and the procedure were the same as reported in Study 3 in Chapter 3 (see § 3.4.1.2).

Results and discussion

To explore the pattern of responses in our non-Masters sample, we conducted a repeated measures ANOVA of the mean prevalence rating, with task (Implied Prevalence vs. Truth-Conditions) and domain (animals vs. people) as between-subjects factors, and property valence (dangerous vs. neutral) as a within-subjects factor. Unlike with the Masters', we did not find a main effect of task ($M_{\text{truth-conditions}}$ = 77.1%; $M_{\text{implied prevalence}}$ = 81.6%), $F(1,97)$ = 2.03, p = .16, ηp^2 = .02. However, like the Masters', we observed a main effect of property valence, $F(1,97) = 16.35$, $p < .001$, ηp^2 = .14, with higher ratings for neutral properties (M = 81.9%) than dangerous properties (M = 76.9%) and an interaction between task and property valence, $F(1, 97) = 7.31$, p =. 008, ηp^2 = .07.

Given this interaction, we examined the simple main effects of property valence within the *Implied Prevalence* and *Truth-Conditions* tasks separately. In the *Implied Prevalence* task, the mean prevalence implied by statements was higher for neutral (M = 85.9%) than dangerous properties (M = 77.4%), $F(1, 97) = 22.86$, $p < .001$, ηp^2 = .19. In contrast, in the *Truth-Conditions* task, the mean prevalence that led participants to accept the statements did

not differ between the two property valences ($M_{neutral}$ = 78.0%; $M_{dangerous}$ = 76.3%), $F(1, 97) = 0.89$, $p = .35$, $\eta p^2 = .009$.

We additionally submitted participants' responses in the *Truth-Conditions* tasks to a logistic mixed-effects model using the glmer command in the lme4 package in R (Bates, 2007). In this model, we included domain (animals = 0; people = 1; between subject); property valence (neutral = 0; dangerous = 1; within subject), and prevalence (.1, .3, .5, .7, .9, 1; within subject) and their interactions as predictors (see Table S3). All predictors were mean centered. We also included *participant* as a random effect. We observed a main effect of prevalence, indicating that generic sentences were more likely to be judged true for higher than lower prevalence levels. However, this main effect needs to be interpreted within the context of a significant property valence by prevalence interaction, which was not observed with the Masters sample (see Figure S2).

Fixed Effects	**Estimate**	***SE***	***p*-value**
(Intercept)	2.39	0.43	< .001
Domain	0.02	0.89	.98
Property Valence	-0.36	0.28	.20
Prevalence	14.02	1.00	< .001
Domain x Property Valence	0.54	0.60	.37
Domain x Prevalence	2.51	1.93	.19
Property Valence x Prevalence	-3.74	1.16	.001
Domain x Property Valence x Prevalence	1.32	2.43	.59
Random Effect		**SD**	
Participant	Intercept	2.76	

Table S3. Logistic regression predicting "true"/ "false" judgments, based on domain, property valence, and prevalence and their interactions in the non-Master sample Study 3.

Post-hoc tests revealed that participants were more likely to endorse generics about dangerous properties than generics about neutral properties at the 10% level (AME = 0.07, SE = 0.02 $p = .004$, 95% CI = 0.02, 0.11) and the 30% level (AME = 0.07, SE = 0.03 $p = .02$, 95% CI = 0.01, 0.13). In contrast, there was no differences between responses to dangerous

and neutral generic predications at higher prevalence levels: 50% level (AME = -0.009, SE = 0.03, p = .77, 95% CI = -0.07, 0.05), 70% level (AME = -0.05, SE = 0.02, p = .049, 95% CI = -0.09, -0.0001), 90% level (AME = -0.02, SE = 0.01, p = .09, 95% CI = -0.04, 0.003), or the 100% level (AME = -0.008, SE = 0.006, p = .17, 95% CI = -0.02, 0.004).

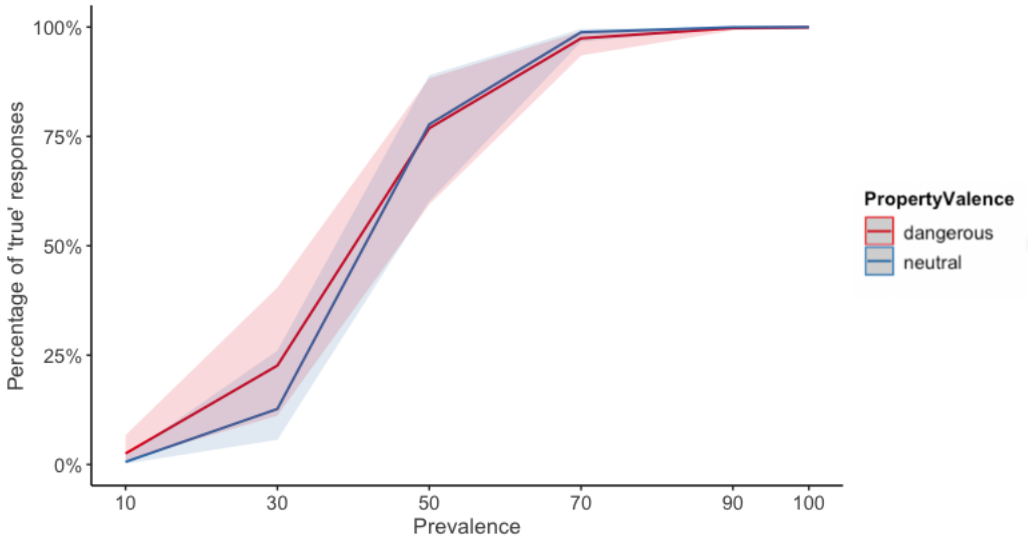


Figure S2. Mean percentage of "true" responses in Study 3 non-Masters sample by property valence and prevalence level.

Unlike with the Masters sample, we did not find an interaction between domain and prevalence. For a complete overview of the mean endorsement percentages at each prevalence level by domain and property valence, see "Truth-Conditions – Study 3 (non-Masters)" in this appendix.

Study 3 - Length differences

Out of an abundance of caution, we conducted a supplementary analysis to ensure that the property valence effects observed in Study 3 were not due to the small length variation of our items. We controlled for predicate length by dropping the two shortest dangerous behavior items (i.e., "attack people" and "kidnap babies") and the two longest neutral

behavior items (i.e., "stamp their feet to great others" and "leave their leftovers on the ground"). The reason for this choice was that dangerous traits and neutral traits had all the same length, whereas dangerous behaviors were slightly shorter than neutral behaviors. Then, we examined participants' responses in the *Implied Prevalence* task based on the remaining 20 items. We again conducted a repeated measures ANOVA with property valence (dangerous vs. neutral) and property content (behavior vs. trait) as within-subjects factors. We found a main effect of property valence, $F(1, 119) = 34.03$, $p < .001$, $\eta_p^2 = .22$, with higher ratings for neutral properties ($M = 92.2\%$) than dangerous properties ($M = 84.4\%$). We additionally found a main effect of property content, with higher ratings for traits ($M = 90.0\%$) than behaviors ($M = 86.6\%$), $F(1, 119) = 18.30$, $p < .001$, $\eta_p^2 = .13$. However, these main effects need to be interpreted within the context of a significant interaction, $F(1, 119) = 13.46$, $p < .001$, $\eta_p^2 = .10$. Given this interaction, we examined the simple main effects of property content for dangerous and neutral properties separately. For the dangerous properties, the mean prevalence implied by statements was higher for traits ($M = 87.5\%$) than behaviors ($M = 81.2\%$), $F(1, 119) = 22.76$, $p < .001$, $\eta_p^2 = .16$. In contrast, for neutral properties, the mean prevalence did not differ between traits ($M = 92.5\%$) and behaviors ($M = 92.0\%$), $F(1, 119) = 0.36$, $p = .55$, $\eta_p^2 = .003$. Since we replicated the same results reported in § 3.4.2.3, we conclude that our findings were not affected by predicate length.

Study 3a

As in Study 3, we recruited a sample of non-Masters for this study. However, when we discovered differences between the pattern of responses across samples (Masters vs. non-Masters) in Study 3, we also wanted to ensure that our sample was consistent across studies. Below we report the results of the non-Masters sample that we collected for this study.

Method

Participants. One-hundred and three adults from the United States (61 men, 42 women; Mean age = 35.44 years; range = 20–68 years) completed the study online and were paid $ 0.75.

Participants were randomly assigned to one of two conditions: *Animals – Truth-Conditions* (*n* = 51) and *People – Truth-Conditions* (*n* = 52). Unlike the participants reported in Chapter 3, these participants had not been granted a Master Worker qualification from MTurk. Participants were 77% White, 10% Black or African American, 7% Latino or Hispanic, 5% Multiethnic, 2% Asian or Asian American. Twenty other participants were tested and excluded from the final sample because they failed the manipulation check (*n* = 6 in the *Animals – Truth-Conditions* and *n* = 14 in the *People – Truth-Conditions*). Four additional participants were tested and excluded for having non-US IP addresses. One participant was tested and excluded from the final sample for having a duplicate IP address. Finally, one participant was tested and excluded from the final sample for not being a native speaker of English.

Materials and procedure. The materials and the procedure were the same as reported in Study 3a in Chapter 3 (see § 3.5.1.2).

Results and discussion

We submitted participants' responses to a logistic mixed-effects model.[3] In this model, we included domain (animals = 0; people = 1; between subject); property valence (neutral = 0; dangerous = 1; within subject), and prevalence (.1, .3, .5, .7, .9, 1; within subject) and their interactions as predictors (see Table S4). All predictors were mean centered. We also included *participant* as a random effect. We observed a main effect of prevalence, indicating that generic sentences were more likely to be judged true for higher than lower prevalence levels. We also observed a main effect of property valence, indicating that statements about dangerous properties were more likely to be accepted than statements about neutral properties. We did not observe this effect with the Masters sample, with whom we found a domain effect instead.

[3] In the pre-registration for this study, we indicated that we would use ANOVAs to analyze the data. However, because the data are binary, it was necessary to analyze the data using non-parametric statistics; we thus opted for logistic regression.

Fixed Effects	**Estimate**	***SE***	***p*-value**
(Intercept)	3.37	0.40	< .001
Domain	0.79	0.74	.29
Property Valence	0.34	0.17	.047
Prevalence	7.50	0.41	< .001
Domain x Property Valence	-0.41	0.34	.23
Domain x Prevalence	1.12	0.78	.15
Property Valence x Prevalence	-0.33	0.54	.54
Domain x Property Valence x Prevalence	-0.66	1.08	.54
Random Effect		**SD**	
Participant	Intercept	3.38	

Table S4. Logistic regression predicting "true"/ "false" judgments, based on domain, property valence, and prevalence and their interactions in non-Master workers in Study 3a.

For a complete overview of the mean endorsement percentages at each prevalence level by domain and property valence, see "Truth-Conditions – Study 3a (non-Masters)" in this appendix.

Truth-Conditions – Study 1

The following is a complete overview of the mean endorsement percentages at the 10-90% prevalence levels in the OTC and ETC tasks of Study 1 reported in this appendix.

		Study 1	
		OTC task	ETC task
Prevalence	10%	35.3%	44.0%
	30%	52.6%	56.0%
	50%	75.0%	79.3%
	70%	81.0%	82.8%
	90%	82.8%	81.9%

Table S5. Mean percentage of "true" responses in the OTC and ETC tasks of Study 1 at the 10-90% prevalence levels.

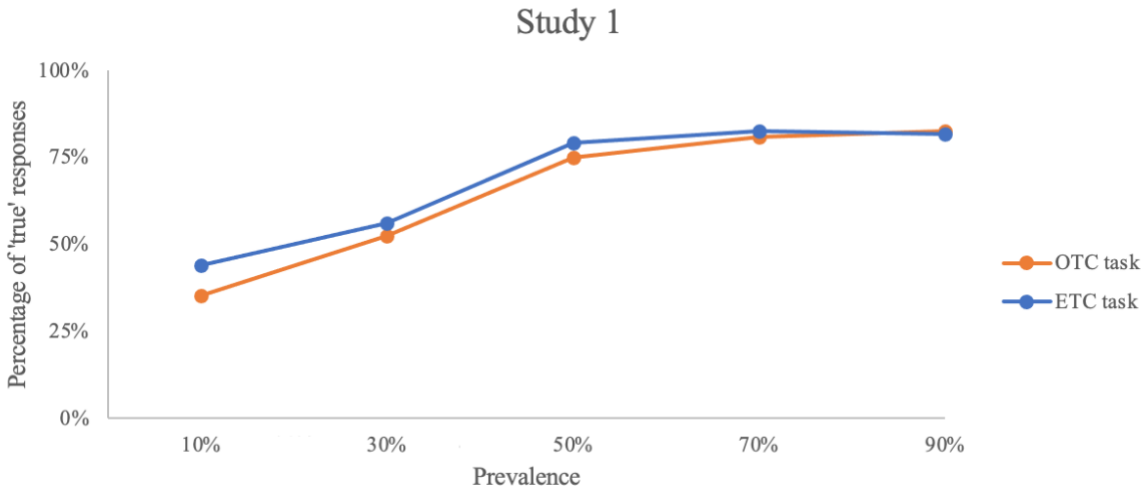


Figure S3. Mean percentage of "true" responses in the OTC and ETC tasks of Study 1 at the 10-90% prevalence levels.

Truth-Conditions – Study 2

The following is a complete overview of the mean endorsement percentages at each prevalence level by domain and property type in the Truth-Conditions task of Study 2 reported in Chapter 3.

Study 2					
		Animals		People	
		Physical	Non-physical	Physical	Non-physical
Prevalence	10%	20.5%	22.3%	38.5%	35.6%
	30%	31.3%	40.2%	42.3%	44.2%
	50%	68.8%	69.6%	61.5%	63.5%
	70%	85.7%	84.8%	72.1%	77.9%
	90%	84.8%	86.6%	79.8%	82.7%
	100%	98.2%	97.3%	96.2%	98.1%

Table S6. Mean percentage of "true" responses in Study 2 at each prevalence level by domain and property type.

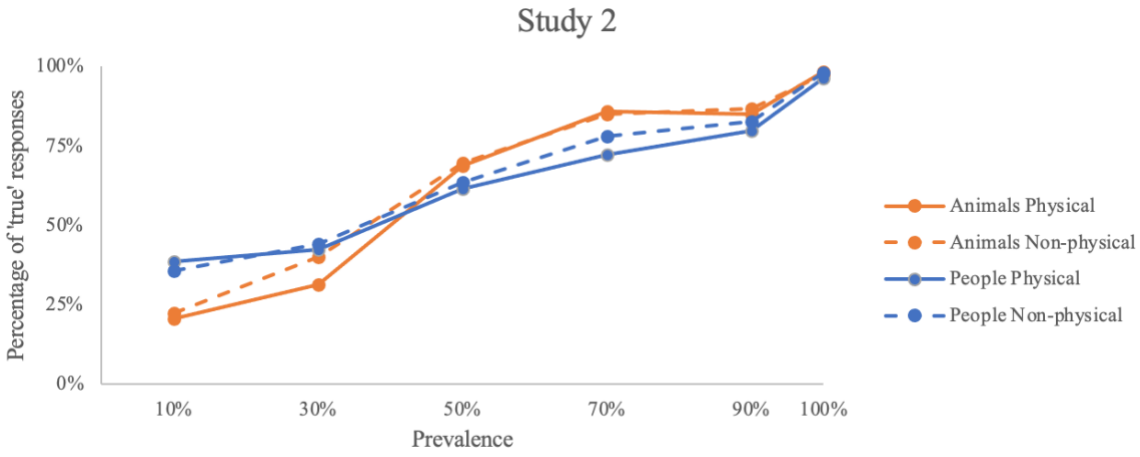


Figure S4. Mean percentage of "true" responses in Study 2 at each prevalence level by domain and property type.

Truth-Conditions – Study 3

The following is a complete overview of the mean endorsement percentages at each prevalence level by domain and property valence in the Truth-Conditions task of Study 3 reported in Chapter 3.

Study 3					
		Animals		People	
		Neutral	Dangerous	Neutral	Dangerous
Prevalence	10%	15.8%	21.1%	29.8%	24.0%
	30%	30.7%	32.5%	32.7%	34.6%
	50%	66.7%	71.9%	72.1%	68.3%
	70%	90.4%	92.1%	89.4%	86.5%
	90%	92.1%	93.0%	88.5%	87.5%
	100%	100.0%	100.0%	94.2%	93.3%

Table S7. Mean percentage of "true" responses in Study 3 at each prevalence level by domain and property valence.

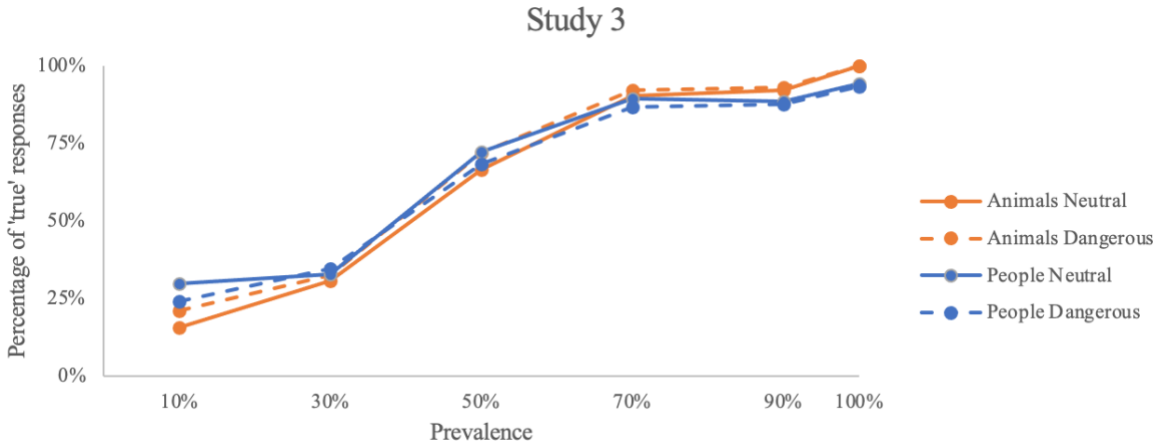


Figure S5. Mean percentage of "true" responses in Study 3 at each prevalence level by domain and property valence.

Truth-Conditions – Study 3a

The following is a complete overview of the mean endorsement percentages at each prevalence level by domain and property valence in the Truth-Conditions task of Study 3a reported in Chapter 3.

Study 3a					
		Animals		People	
		Neutral	Dangerous	Neutral	Dangerous
Prevalence	10%	35.3%	44.0%	55.7%	59.4%
	30%	52.6%	56.0%	67.0%	71.7%
	50%	75.0%	79.3%	87.7%	86.8%
	70%	81.0%	82.8%	89.6%	91.5%
	90%	82.8%	81.9%	94.3%	94.3%
	100%	100.0%	97.4%	99.1%	100.0%

Table S8. Mean percentage of "true" responses in Study 3a at each prevalence level by domain and property valence.

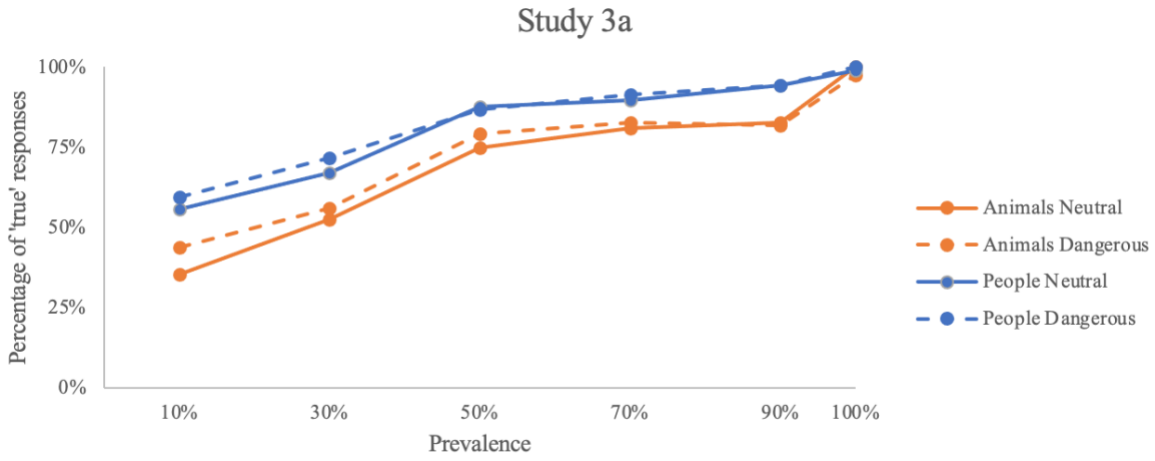


Figure S6. Mean percentage of "true" responses in Study 3a at each prevalence level by domain and property valence.

Truth-Conditions – Study 3 (non-Masters)

The following is a complete overview of the mean endorsement percentages at each prevalence level by domain and property valence in the Truth-Conditions task of Study 3 non-Masters sample reported in this appendix.

Study 3 (non-Masters)					
		Animals		People	
		Neutral	Dangerous	Neutral	Dangerous
Prevalence	10%	10.0%	18.6%	5.6%	13.9%
	30%	21.4%	25.7%	27.8%	27.8%
	50%	70.0%	65.7%	52.8%	61.1%
	70%	94.3%	94.3%	94.4%	91.7%
	90%	98.6%	97.1%	100.0%	94.4%
	100%	100.0%	95.7%	97.2%	100.0%

Table S9. Mean percentage of "true" responses in Study 3 non-Masters sample at each prevalence level by domain and property valence.

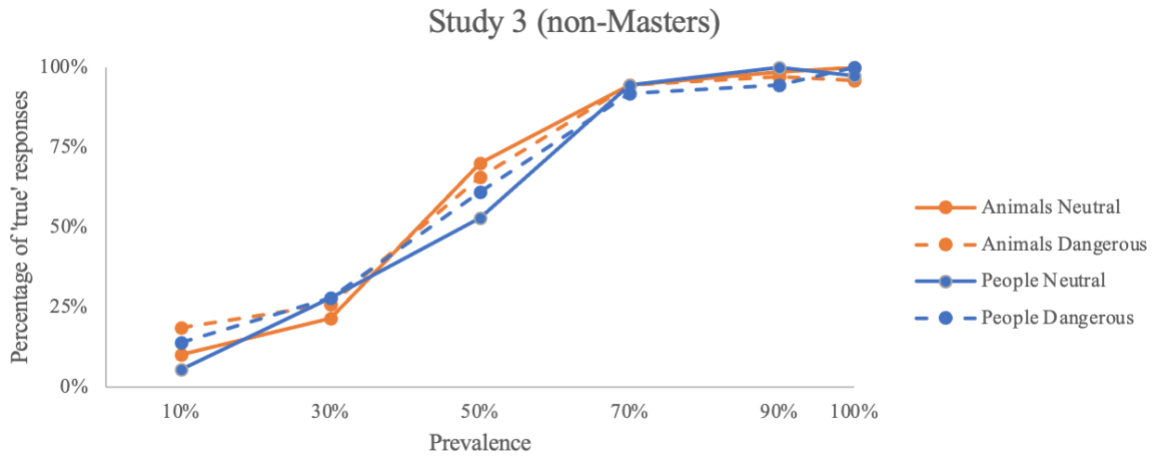


Figure S7. Mean percentage of "true" responses in Study 3 non-Masters sample at each prevalence level by domain and property valence.

Truth-Conditions – Study 3a (non-Masters)

The following is a complete overview of the mean endorsement percentages at each prevalence level by domain and property valence in the Truth-Conditions task of Study 3a non-Masters sample reported in this appendix.

Study 3a (non-Masters)					
		Animals		People	
		Neutral	Dangerous	Neutral	Dangerous
Prevalence	10%	35.3%	42.2%	48.1%	51.9%
	30%	50.0%	57.8%	60.6%	64.4%
	50%	78.4%	81.4%	76.0%	76.0%
	70%	88.2%	93.1%	85.6%	85.6%
	90%	93.1%	95.1%	86.5%	87.5%
	100%	93.1%	95.1%	100.0%	99.0%

Table S10. Mean percentage of "true" responses in Study 3a non-Masters sample at each prevalence level by domain and property valence.

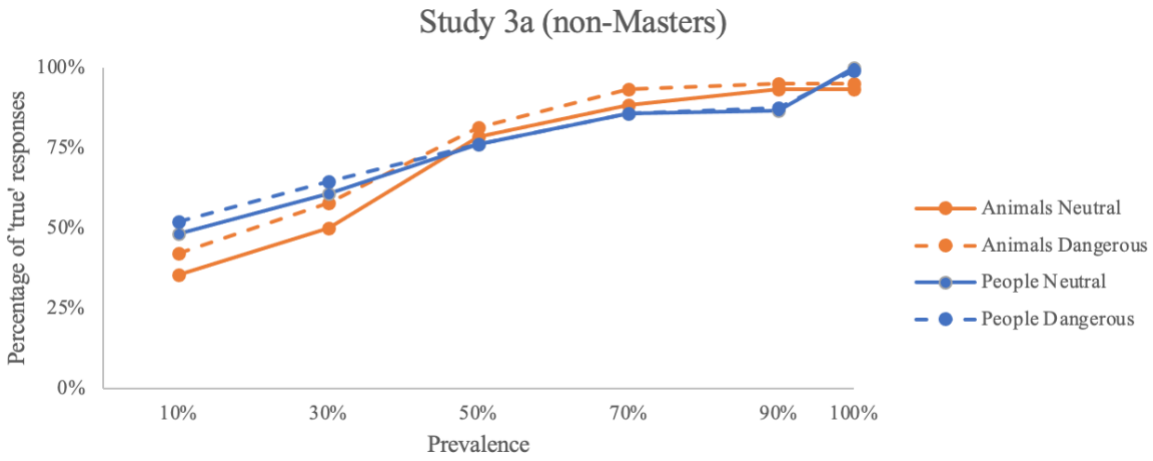


Figure S8. Mean percentage of "true" responses in Study 3a non-Masters sample at each prevalence level by domain and property valence.

Appendix S1

Items from Cimpian et al. (2010)	MORSETHS have silver fur BLINS have red scales ZORBS have orange tails DAITHS have gold stripes MOXES have green shells LUDINOS have yellow legs ELLEPS have pink ears LORCHES have purple feathers GLIPPETS have copper spots THUPS have blue teeth
New items	FRAMS have gray antennae JOTES have brown wings

Items used in Study 1

Appendix S2

Labels		Physical properties	Non-physical properties
MORSETHS	HUXANS	have silver hair/fur	sleep under trees
FRAMS	BEMES	have gray skin	stamp their feet to greet others
LUZAKS	HEABS	have twelve toes	bury their leftover food
JOTES	NOJAS	have red eyes	live underground
KERNS	OGETS	have brown tongues	run in a zigzag pattern
NARES	TEMBAS	have long eyelashes	swim when it is raining
STADES	UPOS	have pointy ears	are nocturnal
TARBS	WHEZAS	have hairy noses	eat insects
ZAVS	WUPTAS	have yellow freckles	dance to communicate
ANEDS	YOXAS	have an extra row of teeth	gather leaves for warmth
EDERS	TOABOS	have large tonsils	are nomadic
GEZOS	VIRDEXES	have short bones	hang upside down

Labels and items used in Study 2

Appendix S3

Neutral properties	Dangerous properties
sleep under trees	carry a deadly virus
stamp their feet to greet others	kill for fun
leave their leftovers on the ground	attack people
hide underground	kidnap babies
run in zigzag patterns	assault for no reason
swim when it is raining	hunt strangers
are nocturnal	are brutal
are short	are dangerous
are messy	are ruthless
are skillful	are bloodthirsty
are nomadic	are violent
are hairy	are vicious

Items used in Study 3

Appendix S4

Animals Condition	
Neutral properties	**Dangerous properties**
have silver fur	have dangerous silver fur. This fur sheds particles that get lodged in your lungs and make it impossible to breathe.
have red scales	have dangerous red scales. These scales secrete a strong venom that kills you on the spot.
have orange tails	have dangerous orange tails. These tails are so long and muscular that they can suffocate you in a matter of minutes.
have gold stripes	have dangerous gold stripes. These stripes deliver a powerful electric shock that's deadly to anyone within a few feet.
have green shells	have dangerous green shells. These shells are so very heavy that they would immediately crush your bones.
have yellow legs	have dangerous yellow legs. These legs are so powerful that a single blow could kill you.
have pink ears	have dangerous pink ears. These ears are home to dangerous parasites that can make you go deaf.
have purple feathers	have dangerous purple feathers. These feathers are as sharp as needles and can easily get lodged in you, causing massive bleeding.
have copper spots	have dangerous copper spots. These spots are home to a contagious fungus that is deadly to anyone who becomes infected with it.
have blue teeth	have dangerous blue teeth. These teeth are razor-sharp and so powerful that a single bite can be lethal.
have gray antennae	have dangerous gray antennae. These antennae are covered with deadly bacteria that attack the body immediately.
have brown wings	have dangerous brown wings. These wings have sharp and long spikes that can easily tear you apart.
People Condition	
Neutral properties	**Dangerous properties**
use silver sprayers	use dangerous silver sprayers. These sprayers shed particles that get lodged in your lungs and make it impossible to breathe.
brew red drinks	brew dangerous red drinks. These drinks are made with a strong venom that kills you on the spot.
braid orange ropes	braid dangerous orange ropes. These ropes are so thick and strong that they can suffocate you in a matter of minutes.
wear gold wristbands	wear dangerous gold wristbands. These wristbands deliver a powerful electric shock that's deadly to anyone within a few feet.
throw green sticks	throw dangerous green sticks. These sticks are so very heavy that they would immediately crush your bones.
wear yellow shoes	wear dangerous yellow shoes. These shoes contain a hidden knife so sharp that a single kick could kill you.
prepare pink ointment	prepare dangerous pink ointment. This ointment contains dangerous parasites that can make you go deaf.
toss purple stones	toss dangerous purple stones. These stones are as sharp as needles and can easily get lodged in you, causing massive bleeding.
use copper dyes	use dangerous copper dyes. These dyes contain a contagious fungus that is deadly to anyone who becomes infected with it.
have blue tools	have dangerous blue tools. These tools are razor-sharp and so powerful that a single stab can be lethal.
build gray fences	build dangerous gray fences. These fences are covered with deadly bacteria that attack the body immediately.
have brown gloves	have dangerous brown gloves. These gloves have sharp and long spikes that can easily tear you apart.

Items used in Study 3a