



Generics and typicality: a bounded rationality approach

Robert van Rooij¹ · Katrin Schulz¹

Published online: 23 July 2019

© The Author(s) 2019

Abstract

Cimpian et al. (2010) observed that we *accept* generic statements of the form ‘ G s are f ’ on relatively weak evidence, but that if we are unfamiliar with group G and we learn a generic statement about it, we still *treat it inferentially* in a much stronger way: (almost) all G s are f . This paper makes use of notions like ‘representativeness’, ‘contingency’ and ‘relative difference’ from (associative learning) psychology to provide a uniform semantics of generics that explains why people *accept* generics based on weak evidence. The spirit of the approach has much in common with Leslie’s cognition-based ideas about generics, but the semantics will be grounded on a strengthening of Cohen’s (1999) relative readings of generic sentences. In contrast to Leslie and Cohen, we propose a uniform semantic analysis of generics. The basic intuition is that a generic of the form ‘ G s are f ’ is true because f is typical for G , which means that f is valuably associated with G . We will make use of Kahneman and Tversky’s Heuristics and Biases approach, according to which people tend to confuse questions about probability with questions about representativeness, to explain pragmatically why people *treat* many generic statements *inferentially* in a much stronger way.

Keywords Generic sentences · Formal semantics · Pragmatics · Associative learning · Probability

The ideas behind this paper were developed in 2016 and first presented at the beginning of 2017 during a conference on lying in Leiden. Afterwards, the ideas were presented by both authors at other occasions in Amsterdam, Berlin, Bochum, Munich, Stockholm, Venice and Zurich. We received many comments after these talks and they all (hopefully) positively influenced our paper. We would like to thank everyone for their comments during these events, in particular Nils Franzén, who acted as a commentator of this paper at the ‘Philosophy Meets Linguistics’ Workshop in Zurich, 2017. We also like to thank Dean McHugh and Shane Steinert-Threlkeld for commenting on the manuscript, including correcting for style and typos. Finally, and most obviously, we would like to thank the reviewers of this paper for their stimulating comments and (mostly) constructive criticism.

✉ Robert van Rooij
R.A.M.vanRooij@uva.nl

Extended author information available on the last page of the article

1 Introduction

Generic sentences come in very different sorts. Consider (1-a) and (1-b).

- (1) a. Tigers are striped.
- b. Mosquitoes carry the West Nile virus.

We take (1-a) to be true, because the vast majority of tigers have stripes. But we take (1-b) to be true as well, even though less than 1% of mosquitoes carry the virus. Most accounts of generics, if they don't stipulate an ambiguity, start from examples like (1-a) and then try to develop a convincing story for examples like (1-b) from here. For our analysis of generics, in contrast, we will take examples like (1-b) as points of departure and then seek to account for more standard examples as well. We will argue that such a road will lead to a more convincing uniform analysis.

Although generics are studied mostly in formal semantics and philosophy, they have recently attracted the attention of cognitive psychologists as well. The reason is that generics play a core role in the way we *learn*, *represent* and *reason* about groups in the world (cf. Leslie 2008). Indeed, generic statements express very basic kinds of inductive generalizations, learned during the process of categorization. A central hypothesis of this paper is that the way we learn new categories is, and remains, of crucial importance for judgements involving those categories. We will argue that generic statements about categories, or groups, express typical information about these groups, and that the way people learn about a group is of crucial importance for what is typical about this group. The notion of *contingency* from associative learning psychology plays an important role in learning, and we will argue that a slight generalization of it is crucial for typicality as well, and thus for the analysis of generics.

In Sect. 2 we will provide a biased overview of some semantic theories of generics. We will concentrate our attention in particular on Cohen's proposal, because that is what our own proposal is built on. After a discussion of theories of categorization in Sect. 3, we will discuss our own semantic account of generics in Sect. 4. According to our own uniform semantic account, a generic like (1-a) is true basically because *relatively many* tigers are striped, except when only tigers are considered in which case the (vast) majority of tigers have to be striped. We will argue that such an account is in accordance with the way we inductively learn categories and how we represent them. This semantic analysis will be closely related to Cohen's treatment of what he calls the 'relative' reading of generics, but such that (under certain circumstances) his 'absolute' reading comes out as a special case. This semantic analysis will give rather weak truth conditions to many generic sentences. In Sect. 5 we will provide a pragmatic explanation of why generics are normally interpreted in a much stronger way, making use of insights of Tversky and Kahneman's (1974) Heuristics and Biases program.

2 Some semantic theories of generics

Generics are sentences that express basic generalities without the use of an explicit quantifier. Generic sentences come in various sorts: they can be expressed using a bare plural, (2-a), a singular indefinite, (2-b), and a definite description, (2-c):

- (2) a. Triangles have three sides.
- b. A triangle has three sides.
- c. The triangle has three sides.

Many interesting observations have been made about the relation between these three types of generic sentences, and about their interpretations. In this paper we will only be concerned with generics of the form (2-a). Some generics of form (2-a), such as (3) are exclusively about kinds, and not about the individuals of that kind:

- (3) T-Rexes are extinct.

We will ignore such examples in this paper as well.

It is often said that generic sentences of the form ‘*Gs are f*’ come in two sorts¹: descriptive ones and non-descriptive, or normative ones. In this section we will be mainly concerned with giving a biased overview of semantic analyses of descriptive generics. But non-descriptive generics will be briefly discussed as well.

Generic sentences are sentences that, by their very nature, express useful generalizations. The main question addressed in the literature is about the type of generalization. First, generic sentences are clearly not *universally quantified* sentences: although not all birds fly (Penguins don’t), (4) is a good generic sentence that most people consider true.

- (4) Birds fly.

Indeed, this is one of the most typical features of generic sentences: they express generalizations that allow for exceptions. But it also need not be the case that almost all, or *most Gs* have feature *f* in order for the generic ‘*Gs are f*’ to be true:

- (5) a. Birds lay eggs.
- b. Goats produce milk.

Although (5-a) and (5-b) are true, it is not the case that the majority of birds or goats have the relevant feature; only the adult female birds and goats do! Moreover, even if most *Gs* are (or are taken to be) *f*, according to Carlson (1977) and others the corresponding generic sentence still doesn’t have to be true, as exemplified by sentences such as the following:

¹ In this paper we will generally use ‘*Gs are f*’ to talk about all generics of the form that we wish to talk about. Intuitively, ‘*G*’ stands for a group and ‘*f*’ for a feature, but one shouldn’t think of the two as having a different type. They both will denote sets, and probabilities or proportions can be given to both in the same way.

- (6) a. Bees are sexually sterile.
 b. Israelis live on the coastal plane.
 c. Books are paperbacks.

According to a natural alternative quantificational proposal, the generic is true exactly if all, or most, *normal*, or *relevant* *Gs* are *f*. There are at least two problems with such an analysis. First, and foremost, without an independent analysis of what it is to be a normal or relevant *G*, such an analysis hardly makes any empirical predictions (cf. Krifka et al. 1995). Second, such an analysis is extensional, and that is taken to give rise to problems exemplified by the following, much discussed, generic (7):

- (7) Mail from Antarctica is handled by Tanja.

This generic can be true, even though we've never gotten any mail from Antarctica. It is normally argued that what such an example points to is a demand for an *intensional* treatment of generics. Arguably, however, (7) is a *normative* generic, and normative generics cannot be given a purely extensional treatment anyway. But, of course, there is a much better reason why generics should not depend on certain actually observed extensionally given sets: if our theory claimed this, we could not account for their *inductive*, or unbounded, character.

According to the *modal nonmonotonic* approach of Asher and Morreau (1995), Pelletier and Asher (1997) and others,² '*Gs* are *f*' is true if and only if for any entity *d* and all worlds in which *d* is a normal *G*, *d* has feature *f*. Such theories want to account for a type of *default instantiation*, that is, for the fact that if all we know is that the sentences '*Gs* are *f*' and '*x* is a *G*' are true, we can normally, or by default, conclude by instantiation that *x* has feature *f*. Proponents of nonmonotonic logic typically argue that what is normal need not have anything to do with proportions. Rather, what is normal is taken to model conventions used in human communication and knowledge organization (cf. McCarthy 1986; Reiter 1987). We don't necessarily want to object to this, although this will probably mean that no uniform explanation can be given for what counts as normal (cf. also Krifka et al. 1995). With Pearl (1988, Ch. 10) one might wonder, then, how useful such a notion of normality really is.³ Moreover, whatever 'normality' is taken to mean, any analysis that wants to account for 'default instantiation' will have problems accounting for the intuition that the following generics are both true.

- (8) a. Ducks lay eggs.
 b. Ducks have colourful feathers.

² Although, for instance, Veltman's (1996) update semantics, which also makes crucial use of the notion of 'normality', was not explicitly used to give a semantics for generic sentences, it was certainly designed to fulfill that role. This holds presumably for many other nonmonotonic logics designed in the last decades of the 20th century as well.

³ Of course, this is reminiscent to a similar criticism one might give to semantics of counterfactuals that make crucial use of a primitive notion of 'similarity'.

In order to predict that (8-a) is true, it must be the female ducks that should be relevant or normal, while it is the opposite sex that is relevant or normal for (8-b).⁴

There is another typical kind of example that is problematic for the analyses discussed so far. Consider the following (seemingly true) generic:

(9) Wolves kill men.

Eckardt (1999) argues that it is very rare for a wolf to encounter human beings, let alone kill them. Hence, in a normal situation, no wolf is in the vicinity of any human being, and hence is not killing any human being. Similarly, a human being normally (at least for the majority of those involved in the academic debate on the meaning of generic sentences) is situated in some European or North-American city or suburb where there hasn't been a wolf for centuries. One could then argue that while normally, a wolf might not be eating a man, it surely must be *disposed* in some way to kill human beings. This is, however, false: normally, wolves actively avoid contact with humans and will often flee when they do encounter a human being. Even those who know this, however, will judge (9) to be true.

Of course, quite a number of examples very similar to (9) are problematic for exactly the same reason.

- (10) a. Mosquitoes carry the West Nile virus.
 b. Ticks carry the Lyme disease.
 c. Sharks attack people.

Sentences (9) and (10-a)–(10-c) are examples of what Leslie (2008) calls 'striking generics'. A generic sentence is striking if only very few of the *Gs* need to have feature *f* for the generic sentence '*Gs* are *f*' to be true. According to Leslie, striking often means 'horrific or appalling'. However, the truth conditions suggested above also seem to hold for familiar examples like (11-a)–(11-c) that are intuitively not 'horrific'.

- (11) a. Frenchmen eat horsemeat. from Schubert and Pelletier (1987)
 b. Bulgarians are good weightlifters.
 c. Dutchmen are good sailors.

Intuitively, a sentence like (11-c) is true not because most Dutchmen are good sailors, but because *relatively many* Dutchmen are. These types of examples motivated Cohen (1999) to claim that generic sentences are in fact *ambiguous*. Generic sentences can both have an *absolute* reading, and a *relative* one. Cohen (1999) develops his theory in terms of probabilities. Cohen believes that generics are objectively true or false. Like most semanticists he believes, for instance, that 'Snakes are slimy' is objectively false, even if most people believe it is true. For this reason Cohen rejects the standard

⁴ Asher and Morreau's (1995) analysis, for instance, predicts that the following principle is valid, if *B* entails *C*: $A > B \models A > C$. Nickel (2009) proposed an account based on 'normality' exactly to deal with these kinds of examples by making the relevant type of normality depend on the predicate being used. We agree with Hoeltje (2017), however, that this analysis is too weak. We won't deal with Nickel's modal analysis at all in this paper, because we have little to add to Hoeltje's criticism.

‘subjectivist’ interpretation of probability, at least for the analysis of generics, and goes for a truly ‘objectivist’ frequency interpretation of probabilities.⁵

His account for generics of the form ‘ G s are f ’ hinges on a notion of contextually supplied alternatives for both f and G . The set of alternatives for feature f , $Alt(f)$, is important for both readings of a generic, while $Alt(G)$ will play a role only for the relative reading. The set $Alt(Cat)$, for instance, will almost always include dogs, and often other pets. Cohen assumes that G and f will be among $Alt(G)$ and $Alt(f)$, respectively. Cohen (1999) proposes the following truth conditions of his absolute and relative readings of generics⁶:

1. An **absolute** generic ‘ G s are f ’ is true if and only if the probability that an arbitrary element of G that has some feature in $Alt(f)$ will have f , is greater than $\frac{1}{2}$. If we assume that all G s have some feature in $Alt(f)$, or limit the set of relevant G s this way, generics are true on this reading iff $P(f/G) > \frac{1}{2}$, or equivalently $P(f/G) > P(\neg f/G)$.⁷
2. A **relative** generic ‘ G s are f ’ is true if and only if the probability that an arbitrary element of G [that has some feature in $Alt(f)$] has feature f is greater than the probability that an arbitrary element of $\bigcup Alt(G)$ [that has some feature in $Alt(f)$] has feature f . Assuming that $\bigcup Alt(G)$ covers the set of all individuals, or if we limit the domain of the probability function P to $\bigcup Alt(G)$, and we assume the same regarding $Alt(f)$, this comes down to $P(f/G) > P(f)$, or equivalently $P(f/G) > P(f/\neg G)$.⁸

Cohen proposes that generic sentences are standardly interpreted in the absolute way, but that sentences that are problematic for many other treatments of generics like (11-a)–(11-c) should be interpreted in the relative way, just as examples like (8-a) ‘Ducks lay eggs’ and ‘Lions have manes’. Presumably, the same is the case for the striking generics (9) and (10-a)–(10-c).

Where Cohen (1999) gives two separate treatments of absolute and relative generics, in Cohen (2001) he provides yet another analysis of *non-descriptive generics* like the following.

⁵ To capture the ‘lawlike’, or ‘unbounded’ character of generics, he makes use of unlimited sequences. According to it, probability statements are not only based on sequences actually observed, but also possible forms such sequences might take in the future as extrapolated from the actually observed ones. As noted by Cohen (1999), objectivity and lawlikeness could be captured as well by a propensity interpretation of probabilities, but he rejects that interpretation on other grounds.

⁶ As noted by Cohen (2001), similar readings have been proposed for sentence involving ‘many’ as well. It is an interesting challenge to explain the clear differences between such sentences and generics, but not one that will be taken up in this paper.

⁷ This equivalence is immediate: $P(f/G) > P(\neg f/G)$ iff $P(f/G) > 1 - P(f/G)$ iff $P(f/G) + P(f/G) > 1$ iff $P(f/G) > \frac{1}{2}$.

⁸ This equivalence is well known, and can be proved straightforwardly:

$$\begin{aligned}
 P(f/G) > P(f) &\text{ iff } \frac{P(f \cap G)}{P(G)} > P(f) \\
 &\text{ iff } P(f/G) > P(G) \times P(f/G) + P(\neg G) \times P(f/\neg G) \\
 &\text{ iff } P(f/G) > \alpha P(f/G) + (1 - \alpha)P(f/\neg G), \text{ with } 0 \leq \alpha \leq 1 \\
 &\text{ iff } (1 - \alpha)P(f/G) > (1 - \alpha)P(f/\neg G) \\
 &\text{ iff } P(f/G) > P(f/\neg G).
 \end{aligned}$$

- (12) a. Bishops move diagonally.
 b. The Speaker of the House succeeds the Vice-President.

In contrast to the generics discussed so far, generics like (12-a)–(12-b) do not describe the world around us. Instead, they seem to express norms or constitutive conventions. Cohen (2001) and following Carlson (1977), proposes that non-descriptive generics should be understood as *rules*. Generic (12-a), for instance, expresses according to Cohen the constitutive rule of what it is to be a bishop. Without a rule like this, there would not even be bishops. Something similar can be said of the other non-descriptive generics like (12-b), although in a less radical way. In contrast to descriptive generics, Cohen (2001) assumes that non-descriptive generics have an underlying logical form that differs radically from their surface form, and that they also have a very different interpretation.

Although we undoubtedly take Cohen’s analysis to be a major step forward compared to other analyses of generics, it is certainly not without problems. The *first problem* for Cohen’s analysis we take to be the claimed *ambiguity*. Can it really be the case that there is no common core between all types of generics? Should absolute generics really be given a separate treatment from relative and non-descriptive ones? Some people have proposed that there exists yet another separate reading of generics: the existential one. One of the appealing features of Cohen’s (1999, 2004a) analysis, we take it, is that existential readings *come out* as a special case of Cohen’s absolute reading of generics, by a specific choice of one of the free variables of his analysis (in this case, *Alt(f)*). For the same reason, we think it is more appealing if (something like) Cohen’s relative and absolute interpretations would come out as special cases of a more uniform treatment of generics than if they are treated as separate readings.⁹ Similarly, the proposal that non-descriptive generics have an underlying logical form that differs radically from their surface form, and that they also have a very different interpretation from descriptive generics is *prime facie*, at least, problematic. Everything else being equal, wouldn’t it be more natural, intuitively, to give all types of generics the same logical form, and have an interpretation of generics that works similarly for them all? Modal accounts of generics that make use of the notion of ‘normality’, like that of Asher and Morreau (1995), have, arguably, a better chance to account for descriptive and non-descriptive generics in a formally uniform way. The reason is that what is normal can be understood both in terms of expectations and in terms of norms. Unfortunately, however, we have seen that such theories (i) require such an extremely context-dependent notion of ‘normality’ that the resulting analysis is hardly insightful and (ii) have problems in particular giving a suitable semantic analysis of what we called ‘relative’ and ‘striking’ generics.

A *related problem* is noted by Leslie et al. (2011), who observe that Cohen’s analysis of relative generics predicts that an example like (13) comes out true.

⁹ We agree with one reviewer that we don’t have a knock-down argument why generics should be given a uniform analysis. Perhaps it is the case that what linguists discuss under the heading of ‘generics’ is simply not a uniform phenomenon. Perhaps. Still, it is common practice in semantics, and taken to be appealing among semanticists, to provide uniform analyses of phrases with (roughly) the same form using contemporary tools. This is what we will do in Sect. 4 of this paper, also informed by psychological work on categorization and learning.

(13) Dogs have three legs.

Example (13) is clearly predicted to be false on its absolute reading. It is predicted to be true on the relative reading, however. The reason is that dogs have a higher probability of staying alive after losing a limb than wolves, foxes, hyenas et cetera, because three-legged dogs will be taken care of by their owners. Furthermore, dogs seem to have a higher probability of losing a limb than, say, hamsters, rabbits, miniature donkeys, and parakeets. Hence, the generic (13) is, we think, falsely predicted to be true on its relative reading relative to both the alternative set Pets and the alternative set Dog-like animals.¹⁰

Of course, Cohen could simply claim that (13) only has an absolute reading, and on that reading the sentence is correctly predicted to be false. But this move only brings us back to the first problem: how should we determine which reading each generic sentence should have?

A *third problem* for Cohen (1999) is related to the use of his *homogeneity condition* to explain away some obvious counterexamples (see Leslie et al. (2011), for example). The counterexamples show that the truth conditions as stated above are far too weak. Sentences (6-a), (6-b) and (6-c) for instance, repeated here as (14-a), (14-b) and (14-c), come out as true on both descriptive readings,¹¹ although the generics are, intuitively, false.

- (14) a. Bees are sexually sterile.
 b. Israelis live on the coastal plane.
 c. Books are paperbacks.

To account for these type of examples, Cohen introduces a *homogeneity condition* for both types of readings of the generic. Rather than just demanding that $P(f/G) > P(-f/G)$ on the absolute reading and $P(f/G) > P(f)$ on the relative one, the above should hold for each cell of a salient partition $\{G_1, \dots, G_n\}$ of G for the absolute reading, and for each cell c_i of the salient partition $\{c_1, \dots, c_n\}$ of $G \cup \bigcup Alt(G)$ on the relative reading. We are sympathetic to this use of partitions: it requires generics to express (inductive) generalization that are *stable*, or invariant. A salient partition of bees into queens (female), workers (female) and drones (male) will correctly predict that (14-a) is false, because neither queens nor drones tend to be

¹⁰ One reviewer doesn't agree, and suggests that veterinarians, who are knowledgeable of the facts, judge (13) to be true. The reviewer bases this on the fact (found after a google search) that veterinarians have a saying 'Dogs have three legs and a spare.' We found this saying as well after a google search, but on this site (<https://nl.pinterest.com/pin/173951604330728121/>). The full quote is

Veterinarians often say that dogs are born with three legs and a spare, a description meant to assure fearful pet owners that life is good for "tripaws" after amputation surgery.

We doubt that this quote provides much motivation for taking veterinarians to accept 'Dogs have three legs' (without the extra 'and a spare').

¹¹ In fact, this is not the exactly the case of (14-a). It is now standardly assumed that honeybee workers are not sterile, but just sexually inactive due to an extreme form of altruism (cf. Seeley 1985). The inhibition of ovary activation is lifted when there is no queen (or larvae). Of course, this would still make the generic 'Bees are sexually inactive' falsely predicted to be true and 'Bees reproduce' falsely predicted to be false. We think we can account for the latter problem by assuming that 'Bees' sometimes can receive a collective interpretation. For the former problem, see below.

sterile. Unfortunately, or so Leslie (2008) complains, it is unclear why (8-a) ‘Ducks lay eggs’ is not predicted to be false for the same reason, due to the salient partition of ducks and animals into male and female ones. But there are several ways to solve this problem, however. For one thing, one could demand that only those partitions are appropriate for the interpretation of ‘Gs are f ’ for which each cell of the partition is compatible with (having) feature f . This is not the case for Leslie’s suggested partition with respect to ‘laying eggs’.¹² Alternatively, Cohen could solve the problem by his assumption that the domain of probability function is restricted by $\bigcup Alt(f)$. Because intuitively $Alt(lay\ eggs) = \{lay\ eggs, give\ live\ birth\}$ and because $\bigcup Alt(lay\ eggs) \approx Females$, Cohen would, or better, could predict that (8-a) is true on its absolute reading.

A fourth potential problem of Cohen’s analysis involves only relative readings: they seem *too weak*, both (i) theoretically and (ii) in terms of predictions. As for (i), in Cohen (1999, 2004b) an explicit motivation is given for the absolute reading of generics in terms of *usefulness*. The absolute reading of a generic is useful because if true, one can make use of it to make default instantiations: if the generic of the form ‘Gs are f ’ is true (on the absolute reading), it means that the chance that an arbitrary G has feature f is pretty high. However, in the same publications no motivation in terms of utility is provided for the relative reading of generics: no motivation is given for why it is useful to know that more G s than alternatives of G s have feature f . And indeed, one might wonder whether relative readings of generics are not too weak to be of any use. As for (ii), Nickel (2012) argues that even for sentences for which the relative reading seems perfectly suited, like for (11-c), Cohen’s analysis is *too weak*.

(11-c) Dutchmen are good sailors.

Consider the case that some Dutch sailors are indeed the best in the world, but that the average Dutch sailor is much worse than the average sailor from other countries. Nickel (2012) claims that in such circumstances (11-c), is not true, although, according to Nickel, Cohen (1999) predicts otherwise. We agree with Nickel’s intuition, but notice that on Cohen’s relative reading (11-c) is actually predicted to be false, if we partition the sailors of each country (by the homogeneity condition) into their good sailors, their mediocre sailors and their bad sailors. Still, we do think that Cohen’s relative readings give rise to empirically too weak readings. As we noted already, on Cohen’s interpretation of relative readings it is falsely predicted that even generics like (13) ‘Dogs have three legs’ come out as true.

There exist other kinds of objections to Cohen’s relative readings as well. According to Leslie (2008), the relative reading of (10-b) ‘Ticks carry the Lyme disease’ gives the wrong truth conditions. She argues that even if most (tokens of) animals of the class to which ticks also belong were mites and many of them carried the Lyme disease, we would—in contrast to Cohen’s prediction on the relative reading—still count a sentence like (10-b) ‘Ticks carry the Lyme disease’ as true. But notice that one way to account for Leslie’s intuition, if correct, would be to count not all *tokens* of objects,

¹² Some readers might notice that this goes against Cohen’s (1999) suggestion to account for (14-c) by appealing to the partition $\{paperbacks, hard\ covers\}$. We will come back to this problem below.

but rather equally many tokens of each relevant *type*. If both mites and ticks were relevant for the relative reading of the generic, we would take into account equally many mites and ticks. This would explain away Leslie's imagined counter-example if other relevant types didn't carry the Lyme disease at all.

Although various aspects of Cohen's analysis of generics have come under attack, in this paper we will build on Cohen's analysis, in particular on his relative readings. Before we give a more thorough defense of (a stronger version of) relative readings of generics based on psychological insights on how we represent and learn categories in the following sections, let us here already point to some, we feel, under-appreciated aspect of such readings. Consider, first, generic sentences that express *comparative relations*, like (15):

(15) Boys are taller than girls.

Although other analyses of generics might be able to account for such readings as well, an analysis in terms of relative readings is almost immediate. The reason is that relative readings are (perhaps implicitly) already treated as comparatives! Let us assume, just for simplicity, an analysis of comparatives as given in Klein (1980): 'John is taller than Sue' is true iff there is a comparison class including (perhaps only) John and Mary such that John is tall with respect to this comparison class, while Sue is not. Similarly, (15) will be true if there is a comparison class c including (perhaps only) boys and girls such that with respect to this comparison class boys are tall and girls are not. Notice that according to the relative reading of the generic sentence 'Boys are tall', the sentence is true iff $P(\text{tall}(c)/\text{boys}) > P(\text{tall}(c)/\neg\text{boys})$. If c consists of only the boys and girls this reduces to $P(\text{tall}(c)/\text{boys}) > P(\text{tall}(c)/\text{girls})$. Similarly, the generic sentence 'Girls are not tall' is true on its relative reading in this context iff $P(\neg\text{tall}(c)/\text{girls}) > P(\neg\text{tall}(c)/\text{boys})$, i.e., iff $1 - P(\text{tall}(c)/\text{girls}) > 1 - P(\text{tall}(c)/\text{boys})$ iff $P(\text{tall}(c)/\text{girls}) < P(\text{tall}(c)/\text{boys})$. As a result, sentence (15) is predicted to be true iff $P(\text{tall}(c)/\text{boys}) > P(\text{tall}(c)/\text{girls})$, which indeed seems to be (almost) the correct result.

Second, there exists a straightforward way to explain why (14-c) 'Books are paperbacks' is a bad generic on its relative reading, without making use of the homogeneity condition. The reason is that the domain of quantification, or the domain of the probability function P , is the union of the alternatives of the feature f = 'are paperbacks'. But most naturally, $\bigcup \text{Alt}(f)$ only contains books, which makes $P(f) = P(f/\bigcup \text{Alt}(f)) = P(f/G)$. As a result, (14-c) is predicted not to be true on its relative reading. Notice that this idea doesn't suffice to make (14-c) false on its absolute reading.

Leslie (2008) observes that although generics are extremely hard to analyse truth-conditionally, we are able to understand and use them successfully with relative ease. She suggests that this is so because generics are the expression of a very primitive default mode of generalizing, which picks up on significant or striking properties and links them to psychologically salient kinds. Indeed, as noted by one reviewer, all languages have generics, even those without number words or other ways of quantification (as claimed, for example, for Pirahã). We completely agree with Leslie's

cognitive approach, and with the idea that the analysis of generics should be closely tied to the way we categorize and make inductive generalizations. It is this insight we want to dwell upon as well. As we will see though, this doesn't necessarily mean that truth-conditional approaches like those of Cohen (1999) are wrong headed. Moreover, or so we will argue, this doesn't mean that generics are as (at least 5-way) ambiguous as Leslie et al. (2011) suggest. Perhaps it is possible to give a more uniform semantic analysis of all types of generic sentences, once we know more about typicality and how we learn inductive generalizations.

3 Typicality and associative learning

People have the natural tendency to classify the objects around them in terms of categories. Objects are grouped together to form a category if they have characteristics in common or are roughly similar to one another. Our thinking in terms of categories reduces the complexity of the world around us considerably. Categorization is one of the most common and most important things we do all the time and crucially influences our behavior. One of the most important functions of categories is that they allow us to make use of induction and generalization. Indeed, the process of categorization itself is perhaps the most basic type of induction, or generalization, we make. It is only natural to assume with Leslie (2008) that generic sentences about categories express these basic generalizations. This suggests that to figure out why we accept certain generic sentences but not others, it is crucial to understand this basic process of categorization. Cohen (2004b) and Leslie in various papers discussed already the relation between categorization and generics. But our focus will be different from either. Cohen (2004b) seeks to motivate his use of homogeneity, which we won't discuss. Leslie concentrates more on learning, but does not discuss standard theories of associative learning, which we will focus on.

3.1 Categorization and typicality

Traditionally, a category was defined in terms of a critical set of attributes the possession of which was taken to be both necessary and sufficient to be a member of the kind. But this traditional conception of categories is now largely abandoned. Typicality plays an important role in more recent theories of categorization and it will play a crucial role in our analysis of generic statements as well. One of the main claims of this paper is that a generic of the form '*G*s are *f*' is true if *f* is a typical feature of *G*s, or that typical members of the category *G* have feature *f*. Typicality is well studied in cognitive psychology. According to prototype theory, groups (or categories) are represented by typical members, rather than by all of them and only them, or by typical features, rather than by necessary and sufficient features, because agents have limited attention and limited recall of examples. But what are a group's typical members or features? According to Rosch (1973), it is the *central*, or *average* members of the group, or the features *most* members have. Centrality is determined in terms of a notion of similarity, which is taken to be based, in one way or another,

on frequency and correlation information. Barsalou (1985) experimentally showed on the basis of a thorough correlational analysis, however, that at least for goal-derived artificial categories, the typical members are instead the category's *ideal* members; those that *best* satisfy the goal. For example, the ideal of the category 'things to eat on a diet' presumably is 'zero calories,' which clearly is not a common, but rather an *extreme* value for members of the category. Idealness can be defined as the extent to which a certain object displays a quality that is directly related to the goal. More recent empirical findings (e.g. Lynch et al. 2000; Palmeri and Nosofsky 2001; Burnett et al. 2005; Ameel and Storms 2006) show that extreme members of a group are also considered typical for many, if not most, other types of categories, namely if categorization is performed in a *contrastive* way. Typical members of a category when defined contrastively have features that *distinguish* them from members of other categories; as such, they highlight, or exaggerate, real differences between groups.¹³

3.2 Associative learning

Typical features for a group, or features that typical members of the group have, are taken to be features that are *representative* for the group. We stated already that our analysis of generics will be based on the process of categorization. A further hypothesis of this paper is that the way we *learn* categories is, and remains, of crucial importance for *judgements* involving those categories.

A popular way to approach the learning of categories involves associative learning based on frequencies and correlations. Much of that psychological research was done before the cognitive revolution in psychology, in classical conditioning. In classical conditioning, what is learned is an association between a cue, *C*, and an outcome, *O*. Pavlov hypothesized that the strength of association between cue and outcome depends on the number of times the two are paired. Subsequent research has revealed, however, that for prediction it is not exactly the number of pairings between cue and outcome that is crucial. In a classic study, Rescorla (1968) showed that rats learn a tone (*C*) → shock (*O*) association if the frequency of shocks immediately after the tone is higher than the frequency of shocks undergone otherwise. Within associative learning psychology, this difference in frequency is known as the *contingency* of the shock on the tone. Rescorla's (1968) central finding was that the higher the contingency of shock on the occurrence of the tone, the more the rats anticipated the fear of a shock. Thus, the higher the contingency, the more useful the tone is as a predictor of the shock. Of crucial importance for our paper is that these experiments show that rats will develop a tone → shock association even if shocks occur only in, say, 12% of the trials in which a tone is present, as long as the frequency of the shocks experienced otherwise is (significantly) lower. Formally, this contingency, or strength of association, between

¹³ We don't assume that some features are *inherently* such that they highlight/exaggerate differences. Instead, whether a feature highlight differences depends on whether members of alternative groups have this feature. The features that exaggerate difference are those that comparatively many members of one group have and the other group(s) don't. They don't have to be special in any other way. Because these are taken to be typical for a whole group, it has the effect that people think that the groups are more different from one another than they actually are, if they think that typical properties are also common properties. (See Sect. 5 for more on this).

C (e.g. tone) and O (e.g. shocks) is measured by $P(O/C) - P(O/-C)$, abbreviated by ΔP_C^O , where P measures frequencies during the learning phase.

Other experiments in the aversive (i.e. fear) and appetitive conditioning paradigms (e.g. Thomas and LaBar 2008) show that the speed of acquisition increases with the *intensity* of the shock. More generally, stronger emotions promote faster learning, more enduring memories, and stronger associations (cf. Chatlosh et al. 1985).¹⁴ One could say that for trained rats, tones play an important role in their categorization of shocks: the tone is a useful predictor and thus provides valuable information to the rat on how to prepare for the future. Moreover, this role of the tone in categorization becomes more entrenched with increased intensity of the shock.

Whereas early work in classical conditioning mostly involved animals, more recent work shows that humans learn associations between the representations of certain cues (properties or features) and outcomes (typically another property or a category prediction) in a very similar way (cf. Gluck and Bower 1988; Cheng and Holyoak 1995; Shanks 1995).¹⁵ For coming to associate feature f with group or category G (via learning), f should be a *distinguishing* feature of G s. However, using contingency, Δ_G^f , to measure this distinguishability is arguably not the most natural measure to account for association. The largest problem is that for contingency the absolute values of $P(f/G)$ and $P(f/-G)$ don't matter, as long as their difference remains the same. In the associative learning literature it is well established, however, that (i) the required difference between $P(f/G)$ and $P(f/-G)$ for learning an association between G and f decreases with an increase of $P(f/-G)$, and (ii) the value of $P(f/G)$ should count for more than the value of $P(f/-G)$: f becomes more strongly associated with G if $P(f/G) = 0.8$ than if $P(f/G) = 0.6$, even if in both cases $\Delta_G^f = 0.1$ (and thus $P(f/-G) = 0.7$ and $P(f/-G) = 0.5$), respectively. A standard way to account for both of these conditions, is to make use of what Shep (1958) calls **relative difference**, which we will denote by $\Delta^* P_G^f$, and which is defined as follows:

$$\bullet \Delta^* P_G^f \stackrel{df}{=} \frac{\Delta P_G^f}{1 - P(f/-G)}.$$

This notion was proposed as well by Cheng (1997) in her analysis of (causal) learning, and Pearl (2000) derives this measure as estimating the probability of (causal) sufficiency, PS. Notice that contingency, and thus distinctness, still plays a major role: for $\Delta^* P_G^f > 0$ it is required that $\Delta P_G^f > 0$, and high ΔP_G^f still results in high $\Delta^* P_G^f$. However, it has the extra effect that for measuring the association between f and G ,

¹⁴ Shanks (1995) shows that ΔP_C^O is the asymptotic value of the weight given to C when the learning task is modeled with a linear associator trained using the Rescorla-Wagner learning rule (Rescorla and Wagner 1972), the most influential learning rule in associative learning which is equivalent to the delta rule used in connectionists models. Interestingly enough, Shanks also points out that associative models using the delta rule learning algorithm can immediately explain the findings of Chatlosh et al. as well that associative learning is influenced by the magnitude or *value* of the outcome. This will be interesting for our final proposal of how to define representativeness later in this section.

¹⁵ Instead of the difference between $P(f/G)$ and $P(f/-G)$ to measure the association, or representativeness, of f with G , others [e.g. Good 1950; Tenenbaum and Griffiths (2001)] use the *quotient* between them. But the idea is very similar: the measure is high if the difference between $P(f/G)$ and $P(f/-G)$ is high.

the value $P(f/G)$ is more important than $P(f/\neg G)$.¹⁶ Moreover, $\Delta^* P_G^f$ increases (with fixed ΔP_G^f), if $P(f/\neg G)$ increases. Thus, $\Delta^* P_G^f$ can be high, although ΔP_G^f is low. For instance, if $P(f/\neg G) = 0.9$, $\Delta^* P_G^f$ will be *ten times* as high as ΔP_G^f (if $\Delta P_G^f > 0$)! Thus, for relatively common features it has the effect that $\Delta^* P_G^f$ will be high, even though ΔP_G^f is relatively low.¹⁷

On the basis of these findings, on our *preliminary* proposal, we measure the **representativeness** of feature f for category or group G by this new notion of relative difference, where $\neg G$ abbreviates $\bigcup Alt(G)$ (and $G \notin Alt(G)$). Because of our use of $\bigcup Alt(G)$ instead of the complement of G , it could well be that $P(f/\bigcup Alt(G))$ is undefined, and thus that also $\Delta^* P_G^f$ is undefined, because $\bigcup Alt(G) = \emptyset$, i.e., if there are no relevant alternatives. We take this to be undesirable, and for that reason assume that $P(B/A)$ is not undefined, if $P(A) = 0$, but instead that $P(B/A) = 0$ under these circumstances.¹⁸ Notice that this has the effect that if $\bigcup Alt(G) = \emptyset$, both ΔP_G^f and $\Delta^* P_G^f$ will come down to $P(f/G)$!

According to the above notion, a representative feature for group G doesn't have to be one that most, or even many, members of the group have. Instead, a representative feature is one that *distinguishes* group G from its alternative(s) (for simplicity denoted by $\neg G$), which is exactly in line with the view on typicality discussed above: those features are representative for a group that highlight, or exaggerate, differences with other groups. Similarly, even though two features f and h are mutually incompatible for members of a certain group (e.g., no peacock both lays eggs and has fantastic blue-green tails), they can still both be representative, because representativeness is

¹⁶ See Cheng (1997) for a proof. Instead of using $\Delta^* P_G^f$, one can also make use of *weighted* contingency to correct for the undesirable consequences of using contingency. Weighted contingency can be defined as follows: $\Delta^\alpha P_G^f \stackrel{df}{=} \alpha P(f/G) - (1 - \alpha)P(f/\neg G)$, with $\alpha \in [\frac{1}{2}, 1]$ (or $\alpha \in \{\frac{1}{2}, 1\}$). In the following section we will use $\Delta^* P_G^f$ to define the truth conditions of generics. But using weighted contingency, one could define the (simplified) truth conditions of generics (for which *Value* is irrelevant) as follows: ' G s are f ' is true iff $\Delta^\alpha P_G^f > \Delta^\alpha P_G^{-f}$. One can show easily that in case $\alpha = 1$, the generic is true iff it is true on Cohen's (1999) absolute reading, while if $\alpha = \frac{1}{2}$, the result is Cohen's relative reading. Yet another proposal would be to make use of a notion of 'weighted relative difference', $\Delta^*_\alpha P_G^f$, defined as follows: $\Delta^*_\alpha P_G^f \stackrel{df}{=} \frac{\alpha P(f/G) - (1-\alpha)P(f/\neg G)}{\alpha - (1-\alpha)P(f/\neg G)}$, with $\alpha \in [\frac{1}{2}, 1]$ (or $\alpha \in \{\frac{1}{2}, 1\}$). If $\alpha = 1$, $\Delta^*_\alpha P_G^f$ reduces to $P(f/G)$, and if $\alpha = \frac{1}{2}$, Shep's 'relative difference' comes out. Though appealing, we won't go for these proposals without an indication what α depends on, for instance because also on our proposal $P(f/G)$ is a special case of $\Delta^* P_G^f$, and the use of the extra parameter α only adds more context-dependence.

¹⁷ Arguably, the increasing effect on our desired notion by an increase in $P(f/\neg G)$ should be higher than it actually is in $\Delta^* P_G^f$. One way to do this is to define a new notion, $\Delta^+ P_G^f$, as follows:

- $\Delta^+ P_G^f \stackrel{df}{=} \frac{\Delta P_G^f}{(1 - P(f/\neg G))^n}$, with $n > 1$ and increasing with higher $P(f/\neg G)$

we won't use this notion in the main text, however.

¹⁸ The standard axiom of probability calculus, $P(A \wedge B) = P(B/A) \times P(A)$ puts no constraint on what the value of $P(B/A)$ should be, in case $P(A) = 0$, because then $P(A \wedge B) = 0$ anyway. If one still finds this unusual analysis of conditional probabilities ad hoc, one can also assume to work in general with the notion of 'weighted relative difference', $\Delta^*_\alpha P_G^f$, as defined in footnote 16, and assume that $\alpha = 1$, if $Alt(G) = \emptyset$, and that $\alpha = \frac{1}{2}$ otherwise. This has exactly the same effect.

calculated at the level of a group, and thus two features that are mutually incompatible for individuals may still be both representative of the group. Only if alternative groups are irrelevant, and categorization is not done in a contrastive way (which according to Ameel and Storms (2006) happens mostly for ‘isolated’ categories), high representativeness reduces to high probability.

It is the frequencies that animals and people were exposed to in the learning phase that count for learning associations. But of course, in many cases people are not exposed to the *actual* frequencies of cues (properties or features) with outcomes (typically another property or a category prediction), but rather with a *distorted* picture of it. Distortion is especially likely to happen when we learn associations through the (social) media. For instance, Kahneman (2011) notes that he had a long-held impression that adultery is more common among politicians than among physicians or lawyers. Only later he realized that this associative belief was probably caused by the fact that the extramarital relations of politicians are much more likely to be reported in the media than the affairs of lawyers and doctors. Still, it is only natural to assume that people will pick up associations from news items in a very similar way that people learn associations through actual exposure. This suggests that learning associations between cues and outcomes from the media also goes via our notions ΔP_G^f or $\Delta^* P_G^f$, but now the frequencies measure not the actual frequencies, but a distorted picture of them via media coverage which is strongly biased towards novelty and poignancy (cf. Kahneman 2011).

Slovic et al. (2004), among others, argue that there exists a deeper link between representativeness of events or features and our emotional reactions to them. Events which give rise to fear and danger come easy to mind not only because of higher media coverage, but also simply because they give rise to strong emotional reactions. Indeed, experiments in operants conditioning, or reinforcement learning, show that learning is, or can be, modified by punishment or reinforcement (cf. Pierce and Cheney 2004). We have seen above that also in this sense, humans are not so different from the animals used in classical conditioning experiments: strong emotions like fear promote faster learning and more enduring memories (cf. Cahill et al. 1995).

To incorporate the insight of Slovic et al. (2004) and of operants-conditioning, we will extend our earlier proposal and propose that representativeness should be defined in a more general way by taking also emotional impact into account. We will measure the **representativeness of f for G** by $\Delta^* P_G^f \times Value(f)$, where $Value(f)$ measures the (absolute value of the) utility/fear/joy of having feature f , or perhaps better, the *intensity* of f . We will abbreviate this measure by ∇P_G^f .

$$\bullet \nabla P_G^f \stackrel{df}{=} \Delta^* P_G^f \times Value(f) = \frac{P(f/G) - P(f/\cup Alt(G))}{1 - P(f/\cup Alt(G))} \times Value(f).$$

We will assume that $Value(f) \geq 1$, and normally that $Value(f) = 1$, meaning that under normal circumstances our notion of representativeness reduces to relative difference, $\nabla P_G^f = \Delta^* P_G^f$. We have seen above that in case $Alt(G) = \emptyset$, then $\Delta^* P_G^f$ reduces to $P(f/G)$. It follows that if $Value(f) = 1$ also ∇P_G^f reduces to $P(f/G)$ under these circumstances.

4 Weak semantics: generics state typicalities

In this section we will claim that a generic of the form ‘Gs are f ’ is true if f is a representative feature for G . Therefore we make the following **semantic claim** (where $f \notin Alt(f)$):

- ‘Gs are f ’ is true iff (i) $\nabla P_G^f > 0$, and (ii) for most $h \in Alt(f) : \nabla P_G^f \gg \nabla P_G^h$ if $Alt(f) \neq \emptyset$ and where \gg means ‘significantly larger’¹⁹

Notice that because of condition (i), a generic sentence ‘Gs are f ’ can only be true if $\Delta P_G^f > 0$, i.e., if the generic is true on Cohen’s relevance reading. Because for our definition of ΔP_G^f , and thus of $\Delta^* P_G^f$, it is important to know what are the alternatives of G , the truth conditions explicitly depend on alternatives. Moreover, they could depend on what the alternatives of feature f are, as well. Notice that if $Alt(f) = \{\neg f\}$ and $Value$ is irrelevant, the generic ‘Gs are f ’ is predicted to be true iff $\Delta^* P_G^f$ is high. The reason is that under these circumstances condition (ii) reduces to $\nabla P_G^f \gg \nabla P_G^{\neg f}$, which basically comes down to the same as ∇P_G^f is high. To see this, notice that for $\Delta P_G^f > \Delta P_G^{\neg f}$ to hold it has to be that $P(f/G) - P(f/\neg G) > P(\neg f/G) - P(\neg f/\neg G)$ iff $P(f/G) - P(f/\neg G) > 1 - P(f/G) - (1 - P(f/\neg G))$ iff $P(f/G) - P(f/\neg G) > P(f/\neg G) - P(f/G)$ iff $P(f/G) > P(f/\neg G)$.

Although making use of alternatives is quite standard in formal semantics and pragmatics in general, and for the analysis of generics in particular, it is hard if not impossible to provide general rules for what the alternative sets should be. Cohen (1999) makes use of the context-dependence of $Alt(f)$ and of $Alt(G)$ (for his relative readings) as well, and our way of thinking of these alternative sets is very similar.²⁰ The alternatives for basic level concepts like ‘dog’, for instance, could, depending on context, be dog-like animals like wolves, foxes, hyenas et cetera, but it could also be other types of pets, like cats, hamsters, rabbits and parakeets. In other cases, however, we might want to think of $Alt(G)$, or $\bigcup Alt(G)$, in an entirely different way. For existential readings of generics, for instance, we might want to think of $\bigcup Alt(G)$ as the set of G s as claimed, or suggested, by an earlier assertion in a conversation. As for $Alt(f)$, for some examples we take it to be just $\{\neg f\}$. But at other times it consists of more specific alternatives. Features like ‘4-legged’ might have ‘ x -legged’ as alternatives, for instance, but this need not be the case: it depends on the interests of the interlocutors in the conversation, and on what is at issue. Similarly, ‘lay eggs’ will most naturally be contrasted with others ways of reproduction, but one can imagine that other features that could distinguish types of animals from one another could be relevant alternatives cases.

We think that it is natural for a generic ‘Gs are f ’ to assume that if $Alt(G) \neq \emptyset$, $Alt(G)$ is such that $G \cup \bigcup Alt(G) \not\subseteq f$. Without this condition, all relevant objects

¹⁹ It is tempting to quantify over *all* alternative features in the definition. But this gives rise to a problem: both ‘Huskies have blue eyes’ and ‘Huskies have upright ears’ probably count as true generics. However, eye color is much more characteristic for this breed of dog, because there are other types of dogs, like the German shepherd, that have pointed ears as well. So, a less demanding quantifier seems more appropriate here. Thanks to one reviewer for bringing up this issue.

²⁰ Although in contrast to Cohen, we will assume that $G \notin Alt(G)$ and $f \notin Alt(f)$. But this assumption is not essential, and is only used to simplify our definitions.

$(G \cup \bigcup Alt(G))$ would have feature f , which would make the generic ‘ G s are f ’ uninformative. For similar reasons we think that it is natural for a generic ‘ G s are f ’ to assume that $Alt(f)$ is such that $\forall h \in Alt(f) : G \cup \bigcup Alt(G) \cap h \neq \emptyset$, i.e., each feature is present in at least some G . As is standard in semantics, we assume for neither type of alternatives that the sets of alternatives be jointly exhaustive. In general, we also won’t assume that the alternatives need to be jointly incompatible.

Let us just explicitly discuss a few examples to illustrate our analysis with straightforward choices of alternatives. Less straightforward choices will be discussed later.

- (16) a. Tigers are striped.
- b. Peacocks have fantastic blue-green tails.
- c. Pit bulls are dangerous.

For (16-a), $Alt(tigers)$ is naturally the set involving big cats like lions, cheetahs, leopards and jaguars, but it might include other types of wild animals as well. Because tigers are the only type of big cats that are striped, and of them (almost) all are striped, $P(striped/tiger) - P(striped/\bigcup Alt(tiger)) \approx 1 - 0$, and thus $\Delta P_{tiger}^{striped} \approx 1$ and $\Delta^* P_{tiger}^{striped} \approx 1$. Assuming that $Value(striped) = 1$, (just like for all alternative features), that is more than enough to make (16-a) true.

For (16-b) it seems natural to compare peacocks with alternative feathered animals, or perhaps with other animals where we typically see them: at a children’s farm. Although only half of the peacocks (the males) have these fantastic blue-green tails (abbreviated as FBGT), this is still much more than what the alternative animals have. Thus, $P(FBGT/peacock) - P(FBGT/\bigcup Alt(peacock)) \approx 0.5 - 0$, and thus $\Delta P_{peacock}^{FBGT} \approx 0.5$ and $\Delta^* P_{peacock}^{FBGT} \approx \frac{0.5}{1-0} = 0.5$. We take it that the alternatives h to feature FBGT are less distinctive for peacocks and thus that $\Delta^* P_{peacock}^{FBGT} \gg \Delta^* P_{peacock}^h$.

The $Value$ of the feature is important for (16-c). Although perhaps only a small percentage of pit bulls are dangerous, say 10%, and although other types of dogs—the natural alternatives to pit bulls—can be dangerous as well, say 2%, the sentence can still be true. The reason is that although $P(dangerous/pit\ bull) - P(dangerous/\bigcup Alt(pit\ bull)) = 0.1 - 0.02$ is only 0.08, and that also $\Delta^* P_{pit\ bulls}^{dangerous} = \frac{0.08}{1-0.02} \approx 0.08$ is small, still $\nabla P_{pit\ bulls}^{dangerous} = \Delta^* P_{pit\ bulls}^{dangerous} \times Value(dangerous)$ is high, because of the large $Value(dangerous)$.

We claim that our general definition of the truth conditions for generic sentences can account for the examples we discussed so far. This is sometimes due to the context dependence of various notions involved. To make that clear, let us first make some general observations concerning some special cases (notice that in contrast to some other proposals, these special cases result not due to ambiguity, assuming the existence of separate readings, but come out as interpretations due to the choice of $Alt(G)$, $Alt(f)$ or $Value(f)$):

1. If $Alt(f) = \{\neg f\}$ and $Value(f) = Value(\neg f)$, a necessary condition for the generic ‘Gs are f ’ to be true is that $\Delta P_G^f > 0$, i.e., $P(f/G) > P(f/\bigcup Alt(G))$, i.e., if the generic is true on Cohen’s **relative reading**. To see this, notice that under the above circumstances, the generic is true iff $\nabla P_G^f \gg \nabla P_G^{\neg f}$. Given the meaning of ∇P , this can only hold if $\Delta P_G^f > \Delta P_G^{\neg f}$.
2. If $\forall h \in Alt(f) : Value(h) = Value(f) = 1$ and $P(f/G)$ is not high, ‘Gs are f ’ is true just in case $P(f/\neg G)$ is very low, and thus f is very **distinctive** for G s.
3. If $\forall h \in Alt(f) : Value(h) = Value(f) = 1$ and $P(f/G)$ is close to 1, ‘Gs are f ’ is true already if $\Delta P_G^f > 0$.²¹ Also, if $Alt(G) = \emptyset$, ‘Gs are f ’ is true iff $P(f/G)$ is high.²² (**standard** generics).
4. If $\forall h \in Alt(f) : Value(h) = Value(f) = 1$, ‘Gs are f ’ holds because $\Delta^* P_G^f = 1$ only if $P(f/G) = 1$ (**definitional**, or **analytic** generics).
5. If ΔP_G^f is only somewhat above 0, and $P(f/G)$ is not high, $Value(f)$ has to be high for ‘Gs are f ’ to be true. (**striking** generics).
6. If $\forall h \in Alt(f) : Value(h) = Value(f) = 1$ and $Alt(G) = \emptyset$ and $Alt(f) = \emptyset$, then ‘Gs are f ’ is true just in case $P(f/G) > 0$, i.e., the **existential** interpretation. The same holds if $Alt(f) = \{\neg f\}$, and $\bigcup Alt(G) \cap f = \emptyset$.

Let us now look at some examples with the above cases in mind.

(1). If $Alt(f) = \{\neg f\}$ and $Value(f) = Value(\neg f)$, the generic ‘Gs are f ’ is true only if $\Delta P_G^f > 0$. Cohen (1999) argues explicitly for this (relative) reading for sentences like ‘Dutchmen are good sailors’ and ‘Bulgarians are good weightlifters’. However, we think, and predict, that in general $\Delta P_G^f > 0$ is not enough for a generic to be true. We claim that this is shown by examples like (13) ‘Dogs have three legs’ which indicated that Cohen’s (1999) relative reading of generics is too weak. Even if $\Delta P_G^f > 0$, f need still not be distinctive enough to characterize the G s. Our analysis demands more than $\Delta P_G^f > 0$, in particular if $Value$ is irrelevant and the relevant feature is uncommon. In those cases we demand that $\Delta P_G^f \gg \Delta P_G^h$ for most alternative h s (which comes down to $\Delta P_G^f \gg \Delta P_G^{\neg f}$, if $Alt(f) = \{\neg f\}$). That is, the generic is true only if there is almost no relevantly salient alternative h that is a more distinctive feature for being a G than f is. This, we take it, won’t in general be the case for ‘three-legged’ with respect to dogs.²³

²¹ There are two cases to consider. (i) $P(f/G)$ is high and $P(f/\bigcup Alt(G))$ is low. Then it holds that ΔP_G^f is high, and thus also $\Delta^* P_G^f$ is high; (ii) both $P(f/G)$ and $P(f/\bigcup Alt(G))$ are very high, but $\Delta P_G^f > 0$. Then it still will be the case that $\Delta^* P_G^f$ is high, because then $1 - P(f/\bigcup Alt(G))$ will be low.

²² We have observed before that if $\bigcup Alt(G) = \emptyset$, $\Delta^* P_G^f = P(f/G)$. Thus, ‘Gs are f ’ is true iff for most $h \in Alt(f) : P(f/G) \gg P(h/G)$.

²³ Although, arguably, $P(3legged/Dog) - P(3legged/\bigcup Alt(Dogs)) > 0$, the value $\Delta P_{Dog}^{3legged}$ will be extremely small. Because $1 - P(3legged/\bigcup Alt(Dogs)) \approx 1$, $\Delta^* P_{Dog}^{3legged}$ too will be extremely small. Notice that although $\Delta^* P_{Dutchmen}^{good\ sailors}$ might be small as well in the 17th century, being a good sailor might have been the most salient feature in terms of which Dutchmen could be set apart.

(2). A generic sentence ‘ G s are f ’ is true if f is *very distinctive* for G s. We claim that a generic like ‘Tigers have stripes’ is considered true because ‘having stripes’ is (among the relative alternative features $Alt(f)$) among the most distinctive features of tigers. A generic sentence like ‘Germans are right-handed’, on the other hand, is not predicted to be true if Germans are contrasted with, say, other European citizens (indicated, for instance, by means of contrastive stress on ‘Germans’), simply because ‘being right-handed’ does not distinguish Germans in any significant way.

Our semantic analysis of generics also explains the following example that is paradoxical to many other theories: although only (adult) male lions have manes, (17-a) is an accepted generic, but (17-b) is not.

- (17) a. Lions have manes.
b. Lions are male.

The reason is that compared to lions, relatively few other animals have manes, but it is not the case that compared to other animals relatively many lions are male. Our analysis thus correctly claims that ‘ G s are f ’ can be true and ‘ G s are h ’ false, although $P(h/G) > P(f/G) < \frac{1}{2}$.

Our analysis of generics is based on typicality and as such is very similar to analyses based on *prototypicality*. But the linguistic literature has not been friendly to such approaches. Let’s see whether we can rebut the troubles typically discussed. First, this approach is accused of for simply passing on the problem of generics to a new problem of what it means to be prototypical. But this can’t be a serious problem anymore, given our very explicit proposal, based on psychological research, for what it means to be (proto)typical, or representative. A second problem normally discussed is that this approach cannot deal with the fact that the following two sentences both seem to be true:

- (18) a. Peacocks have fantastic blue-green tails.
b. Peacocks lay eggs.

The reason why this is taken to be a problem is that the proposal to handle generics in terms of prototypicality is mostly taken to be that the sentence ‘ G s are f ’ is true just in case the prototypical G s have feature f . Hence: ‘Tigers have stripes’ is true if and only if all (proto)typical tigers have stripes, and ‘Turtles live to be over 100 years old’ if and only if this is true of the most typical turtles of all. Natural as such an analysis might be, it falsely predicts that (18-a) and (18-b) cannot both be true, because it is not the case that the typical peacock both has a blue-green tail and lays eggs, simply because there is no peacock that is male *and* female. Fortunately, our analysis differs from the one that is criticized. On our analysis it is possible that ‘ G s are f ’ and ‘ G s are h ’ are true—because both $\Delta^*P_G^f$ and $\Delta^*P_G^h$ are high—even though f and h are, in fact, incompatible. This, obviously, is well possible: compared to other animals (in general), many peacocks have beautiful blue-green tails and many peacocks lay eggs.²⁴ (Though we have discussed in Sect. 2 another reason why (18-b), like (8-a) ‘Ducks lay eggs’, only ranges over female animals.) What is predicted not

²⁴ Leslie (2008) provides a stronger argument against the prototype theory. She argues that not only (18-a) and (18-b) are true, but also ‘Peacocks have fantastic blue-green tails and lay eggs’. Perhaps a more

to be possible is that both ‘ G s are f ’ and ‘ G s are $\neg f$ ’ are true (if $\neg f \in Alt(f)$ and $f \in Alt(\neg f)$), which is as it should be according to Hoeltje (2017).

(3). Although f being very distinctive for G s is predicted to be a *sufficient* condition for generics of the form ‘ G s are f ’ to be true, it can’t be a *necessary* condition. This is clearly shown by examples like (19-a)–(19-d), which are all undoubtedly good:

- (19) a. Humans are mortal.
 b. Birds are warm blooded.
 c. Dogs are 4-legged.
 d. Lions are mammals.

Intuitively, these generics are true simply because (almost) all of the mentioned animals have the relevant features, i.e., when the generic is true on Cohen’s (1999) absolute reading. Our analysis can account for such cases as well. Notice, first, that although in all the above cases having the feature f hardly distinguishes the animals involved, G s, from their alternatives, $\bigcup Alt(G)$, it is arguably still the case that $P(f/G) > P(f/\bigcup Alt(G))$ for natural choices of $Alt(G)$. This is perhaps obvious for (19-b)–(19-d), but makes sense (also to assure that not all objects taken into account are presupposed to have feature f) even for (19-a) by taking some immortals into account. As a result, $\Delta^* P_G^f > \Delta^* P_G^{\neg f}$.

But on our analysis, this is not enough for a generic of the form ‘ G s are f ’ to be true, if $Value(f)$ is irrelevant (meaning that $\forall h \in Alt(f) : Value(h) = Value(f) = 1$). Our analysis requires that $\Delta^* P_G^f \gg \Delta^* P_G^{\neg f}$, if $Alt(f) = \{\neg f\}$. Fortunately, this requirement is, arguably, met also for examples like (19-a)–(19-d). The reason is that the features involved in these sentences are rather common among all animals, which makes $P(f/\bigcup Alt(G)) \approx 1$ (though not quite). As a result, although ΔP_G^f is small (but positive), the value of $\Delta^* P_G^f = \frac{\Delta P_G^f}{1 - P(f/\bigcup Alt(G))}$ is still high, because $1 - P(f/\bigcup Alt(G))$ is low (but > 0). For instance, if $P(f/G) = 1$ and $P(f/\bigcup Alt(G)) = 0.95$, then $\Delta P_G^f = 1 - 0.95 = 0.05$ is small, but $\Delta^* P_G^f$ still receives its maximal value, i.e., $\Delta^* P_G^f = \frac{0.05}{0.05} = 1$. Thus, our analysis predicts that some features can be representative for a group, even though they are not very distinctive for the group, because the features are rather common. Thus, on our analysis distinctiveness of a feature f for group G is a sufficient but not necessary condition for the corresponding generic to be true.

What does our analysis predict, if it is clear that in the interpretation of (19-d), for instance, we are only considering lions? This can be the case because (19-d) is given as answer to the question ‘What kind of animals are lions?’ We have seen above that

Footnote 24 continued

convincing example is given by Nickel (2009): ‘Elephants live in Africa and Asia’ (although according to one reviewer this is unproblematic, because it involves a commulative reading). Note that on our analysis it might well be possible that for two mutually incompatible features like f and h it could be that $\neg \exists g \in Alt(f) : \Delta^* P_G^g \gg \Delta^* P_G^f$ and $\neg \exists g \in Alt(h) : \Delta^* P_G^g \gg \Delta^* P_G^h$, even if $Alt(f) = Alt(h)$. What is obviously not possible on our analysis is that for the *conjoined* feature $f \wedge h$ it holds that $\Delta^* P_G^{f \wedge h}$ is high. Thus, for such cases, ‘ \wedge ’ must have wide scope (on a non-commulative reading, if such a reading exists).

because of our assumption that $P(B/A) = 0$, if $P(A) = 0$, it immediately follows that $\Delta^* P_G^f$ reduces to $P(f/G)$ in case $\bigcup Alt(G) = \emptyset$ (and otherwise stays as before). Thus, if only lions are considered, the generic ‘Lions are mammals’ is on this modification predicted to be true iff $P(f/G) \gg P(\neg f/G)$ (if $Alt(f) = \{\neg f\}$), like Cohen (1999) requires for his absolute reading. Notice, though, that on our proposal we can account for Cohen’s absolute reading without stipulating that it is a separate reading: it just falls out as the interpretation because $\bigcup Alt(G) = \emptyset$.²⁵

(4). The use of relative difference, $\Delta^* P_G^f$, instead of contingency, ΔP_G^f , is arguably useful to account for **mathematical** examples. For instance, we have to explain why (2-a), repeated here as (20-a), is a good generic, but (20-b) is not, although the vast majority of prime numbers are odd.

- (20) a. Triangles have three sides.
- b. Prime numbers are odd.

A natural way to account for mathematical generics in terms of our framework is to demand that $\Delta^* P_G^f$ has the *maximal* value, i.e. $\Delta^* P_G^f = \frac{P(f/G) - P(f/\bigcup Alt(G))}{1 - P(f/\bigcup Alt(G))} = 1$. It is easy to see that this is the case only if $P(f/G) = 1$. This give rise to the correct prediction that (20-a) is true, but that (20-b) is false, just as it should be.

Let us now consider **non-descriptive generics** like (12-a) ‘Bishops move diagonally’. At least since Kripke (1972/1980) we know that identity statements can be used in two different ways: (i) to state the identity of meaning (intension) of the two terms, or (ii) to fix, or define, the meaning of one term in terms of the meaning of the other. Generic sentences are very much like identity statements in this sense, and can be used in those two similar ways. Kripke explains the *a priori* character of a sentence like ‘Stick *S* is one meter long’ when talking about the ideal stick, or standard meter, preserved in Paris ever since the French Revolution by the second use of identity statements. Stalnaker (1978) proposed to account for this making use of the idea that worlds not only determine what is the extension of a term, but also its intension (or meaning). This allows him to distinguish two kinds of propositions that can be expressed by sentence *A* in world *w*: (i) the standard proposition $\{v \in W : A \text{ is true in } v, \text{ given the meaning of } A \text{ in } w\}$, and (ii) the diagonal proposition $\{v \in W : A \text{ is true in } v, \text{ given the meaning of } A \text{ in } v\}$. To account for the *a priori* character of identity statements, Stalnaker uses the diagonal proposition expressed by such a statement.

We propose that on a definitional use of a generic of the form ‘*G*s are *f*’, having *f* is a *necessary* condition for being a *G*. Thus, it is required that $P(f/G) = 1$. We have seen above that for $\Delta^* P_G^f$ to be 1, it has to be the case that $P(f/G) = 1$. Thus, it seems natural for definitional uses of generics to demand that $\Delta^* P_G^f = 1$. To account for the

²⁵ If one doesn’t like the alternative treatment of conditional probability, one could also make use of the notion of ‘weighted relative difference’, $\Delta_\alpha^* P_G^f$, as defined in footnote 16. Alternatively, $\bigcup Alt(G)$ could perhaps be thought of as the set all *G*s *imaginable*, i.e., the *G*s in all (im)possible worlds. In case we can imagine that lions are not mammals, the value of $\Delta^* P_{Lion}^{mammal} = 1$, and the sentence is predicted to be true. This move is reminiscent to, and can (we think) be motivated by, Fernando and Kamp’s (1996) similar move for the analysis of ‘many’.

definitional or *a priori* character of definitional used generics, we follow the basic idea of Krifka (2013) and concentrate on the diagonal proposition expressed by the generic.

(5). Next, if $\Delta P_G^f > 0$ but small, and $P(f/G)$ is not high, $Value(f)$ has to be high for ‘Gs are f ’ to be true. Recall that $Value$ was brought in to take over some insights from fear-conditioning. We claim that it is exactly this that makes our analysis immediately account for **striking generics** like (10-b) ‘Ticks carry the Lyme disease’ and (21),

(21) Sharks attack swimmers.

which are problematic for default-based approaches (e.g. Asher and Morreau 1995). Indeed, Leslie (2008) notes that ‘striking’ often means ‘horrific or appalling’, which means having a high $Value$.

What about examples like (11-a), repeated here as (22)?

(22) Frenchmen eat horsemeat.

Here, fear doesn’t seem to play any role. Still, Schubert and Pelletier (1987: 338) point out that, for North Americans, eating horsemeat can be thought of as striking, because by many taken to be ‘disgusting’: very uncommon and giving rise to non-neutral feelings. Learning that members of a group have a ‘disgusting’ habit more than on average can be very interesting. Making use of Shannon’s Information Theory, we can at least account for the ‘very uncommon’ aspect. f is uncommon only if f ’s informativity, $\text{inf}(f)$, is high. The latter notion is defined as $\log_2 \frac{1}{P(f)} = -\log_2 P(f)$. According to this definition, $\text{inf}(f)$ receives a high value exactly if $P(f)$ is small. ∇P_G^f now comes down to $\Delta^* P_G^f \times \text{inf}(f)$,²⁶ if f gives rise to non-neutral feelings. If f is a very common feature, $\text{inf}(f)$ will be small and ∇P_G^f can be high only if $\Delta^{(*)} P_G^f$ is high. For very uncommon features, however, ∇P_G^f can be high even if ΔP_G^f and $\Delta^* P_G^f$ are low (but > 0). We claim that this is going on for generics like (22).

Are the predicted readings strong enough to make generics that we predict true also worthwhile to say? Notice that utility, or value, is a crucial part of our analysis: even though only a very small percentage of ticks actually carry the Lyme disease, it is still worthwhile to find out that (10-b) is true. Although there is only a small probability that the tick you see carries the disease, the disease is serious enough to take measures: the expected value of taking measures is higher than of doing nothing in case you see an arbitrary tick. This suggests that the usefulness of a striking generic like (10-b) is very similar to the usefulness of a standard generic like ‘Tigers have stripes’. Both can be explained in terms of default instantiation, meaning that for both types of generics of the form ‘Gs are f ’, we take it that, or act upon as if, any individual G has feature f . But whereas the default conclusion makes sense for standard generics for reasons of expectation, or probability (the probability that this individual G has feature f is high), for striking generics the default should be explained in terms of the more general

²⁶ What $\text{inf}(f)$ is meant to measure is the informativeness that an arbitrary x has feature f . It is well known that informativity can be thought of as a special case of $Value$, or utility. Bernardo (1979) even shows under which circumstances it reduces to it. For more discussion, see, e.g. van Rooij (2004).

notion of expected value of possible actions. A sentence like ‘Pit bulls are dangerous dogs’, for instance, is a good generic, or so we propose, because if knowing that a dog is a pit bull increases (from, say, 2 to 10%) our expectation that the dog will be dangerous and thus increases our ability to evaluate possible actions accordingly.²⁷ Similarly, it is useful for categorization to concentrate on distinguishing features.

Value doesn’t have to have anything to do with survival. Often enough it is just well-being that counts. Consider (22) again. It is apparently worthwhile to say, or learn, that this sentence is true, although we would survive equally well without this knowledge. But then, notice also that our survival value of learning that Madonna has a new boy toy is practically zero: it wouldn’t change most of our behavior in any way. Still, the media finds it interesting enough to report about it. Similarly, even if there is no real reason to be bothered by what relatively many Frenchmen tend to eat, it can still increase our well-being to gossip about about these Frenchmen by discussing their strange disgusting eating habits. (Recall that according to Dunbar (1996), strengthening social bonds is at least one of the main functions of language use, and gossip helps a lot).

(6). Although for ‘Gs are f’ to be true, normally a significant portion of Gs have to have feature f, sometimes the generic seems to have a much weaker *existential* interpretation (cf. Lawler 1973; Krifka et al. 1995). Indeed, existential generics like (23) seem to pose a problem for nearly any analysis of generics.

- (23) A. Indians don’t eat beef.
- B. No! Indians [do]_F eat beef.

Cohen (1999, 2004a), however, is able to account for existential readings of generics by assuming that these are interpreted on his absolute reading with $Alt(f) = \{f\}$, because then $P(f/G \cap \bigcup Alt(f)) > \frac{1}{2}$ just in case $P(f/G) > 0$. On our proposal we get the same result if we assume that $Alt(f) = \emptyset$ and $Alt(G) = \emptyset$ and if we limit the domain of the probability function to $f \cup \bigcup Alt(f)$ (recall that in contrast to Cohen (1999) we assumed that $f \notin Alt(f)$). Although formally appealing, the proposal is not the only one that accounts for existential interpretations. We can formally account for existential interpretations as well by assuming that *Value* is irrelevant, $Alt(f) = \{\neg f\}$ and thinking of $Alt(G)$ such that $P(f/\bigcup Alt(G)) \approx 0$ (but > 0).²⁸

²⁷ For other cases where expected utility, instead of just expectation, is claimed to be relevant in natural language semantics, see, among others, van Rooij (2003) and Malamud (2012).

²⁸ To show this, notice that $\Delta^* P_G^f > \Delta^* P_G^{\neg f}$ only if $[P(f/G) - P(f/\bigcup Alt(G))] > [P(\neg f/G) - P(\neg f/\bigcup Alt(G))]$. By taking $P(f/G) = p$ and denoting by \hat{x} a number that is very close to x , what follows is that $[P(f/G) - P(f/\bigcup Alt(G))] = [p - \hat{0}] > [(1-p) - \hat{1}] = [P(\neg f/G) - P(\neg f/\bigcup Alt(G))]$. This reduces to $\hat{p} > -\hat{p}$, which holds exactly if $p > 0$, i.e., if $P(f/G) > 0$: the existential interpretation. But how is this compatible with the demand that $\Delta^* P_G^f \gg \Delta^* P_G^{\neg f}$ has to hold for the generic to be true? It is, because what $[p - \hat{0}] = \hat{p} > -\hat{p} = (1-p) - \hat{1}$ results from is the idea that for the generic to be true it *at least* has to be the case that $\Delta P_G^f > \Delta P_G^{\neg f}$. What we demanded however, involves $\Delta^* P_G^f$ and $\Delta^* P_G^{\neg f}$: it should be that $\frac{\Delta P_G^f}{1 - P(f/\bigcup Alt(G))} \gg \frac{\Delta P_G^{\neg f}}{1 - P(\neg f/\bigcup Alt(G))}$. Although $\Delta^* P_G^f$ is close to $\Delta P_G^f = \hat{p}$ because $P(f/\bigcup Alt(G)) \approx 0$, $\Delta^* P_G^{\neg f}$ is approximating $-\infty$ and thus much lower than $\Delta P_G^{\neg f} = -\hat{p}$, because $P(\neg f/\bigcup Alt(G))$ is approximating 1, and thus $1 - P(\neg f/\bigcup Alt(G))$ approximates 0.

We don't know which of these two ways to deal with existential interpretations is conceptually most appealing, and whether we can account for the range of generic examples (e.g. 'Even Indians eat beef') discussed by Cohen (2004a) that receive an existential interpretation on the second way remains to be seen. Notice that on our second analysis we don't assume that $\bigcup Alt(f)$ limits the domain over which the probability function ranges. This, assumption, however, is crucial for Cohen (2004a), and much of his arguments for why $Alt(f) = \{f\}$ in these relevant cases seem to depend on that. Moreover, we feel it is natural that in case of existential interpretations $\bigcup Alt(G) \subseteq G$. For (23), for instance, we think it is natural to think of $\bigcup Alt(G)$ as the set of Indians that verify what is said by A: the Indians that don't eat beef. The result is that B claims that A did not take all Indians into account, not the ones that do, in fact, eat beef. Arguably, something similar holds for 'Even Indians eat beef'.

Before we leave this section dealing with the semantic analysis of generics, let us finally contrast our analysis briefly with some alternative analyses we have not mentioned yet, mainly to illustrate what is special about our own analysis.

We have proposed that generic sentences should be analyzed in terms of representativeness, and that the representativeness of feature f for group G should be measured by ∇P_G^f , which is defined making use of contingency, ΔP_G^f . Notice, now, that ΔP_G^f is closely related to two other interesting measures, $\frac{P(f/G)}{P(f)}$ and $\frac{P(f/G)}{P(f/-G)}$.²⁹ It is remarkable that $\frac{P(f/G)}{P(f)}$ and $\frac{P(f/G)}{P(f/-G)}$ have been proposed as measures of *stereotypicality* of f for G within social psychology (McCauley et al. 1980; Schneider 2004) and economics (Bordalo et al. 2016), respectively. Indeed, just like ΔP_G^f (and $\Delta^* P_G^f$), also $\frac{P(f/G)}{P(f)}$ and $\frac{P(f/G)}{P(f/-G)}$ give those features a high value that are distinctive for group G , and thus highlight or exaggerate differences between groups. From these proposed measures of stereotypicality, together with our own proposal involving representativeness, it naturally follows that in case $\Delta^{(*)} P_G^f$ is high, we could say that f is a stereotypical feature of G . Is that already enough evidence to propose that ' G s are f ' is a good and true generic?

Indeed, a number of authors, including Declerck (1996) and Geurts (1985), have proposed that generics are about stereotypical properties. This account has been criticized by Krifka et al. (1995), however, and is largely abandoned in the literature. But let us see whether the arguments used by these authors are valid for our interpretation of stereotypicality. A first argument used by Krifka et al. (1995) is that stereotypes

²⁹ Contingency, ΔP_G^f , is closely related to the standard notion of relevance, $P(f/fG) - P(f)$:

$$\begin{aligned} P(f/G) - P(f) &= P(f/G) - [(P(f/G) \times P(G)) + (P(f/-G) \times P(-G))] \\ &= P(f/G) - [\alpha P(f/G) + (1 - \alpha)P(f/-G)], \text{ with } 0 \leq \alpha \leq 1 \\ &= (1 - \alpha)P(f/G) - (1 - \alpha)P(f/-G) \\ &= \beta [P(f/G) - P(f/-G)], \text{ with } 0 \leq \beta \leq 1. \end{aligned}$$

Obviously, $\frac{P(f/G)}{P(f)}$ behaves monotone increasingly with $P(f/G) - P(f)$, just as $\frac{P(f/G)}{P(f/-G)}$ behaves monotone increasingly with $P(f/G) - P(f/-G)$. Deciphering the German Enigma Code, Turing once introduced the measure $\log \frac{P(f/G)}{P(f/-G)}$, which behaves monotone with these latter measures as well, given the nature of logarithmic functions. Although these notions are not all ordinal equivalent, all these measures are high if $P(f/G) - P(f)$ is high.

are just widely acknowledged *beliefs* within a speech community, while the truth of a generic depends on actual *facts*: even if uttered in a culture where everybody believes that snakes are slimy, they argue, ‘Snakes are slimy’ is still false. This argument is obviously invalid with respect to our analysis of stereotypes, however, if we base our analysis not on a subjective probability function, but on objective frequencies, or propensities. The truth of a generic is then predicted to depend on actual facts. A second counterargument of Krifka et al. (1995) is that stereotypes are tied to well-known groups or kinds, while generics are often not about any of those things. But, again, we don’t see why this could be problematic for our analysis: for the calculation of $\Delta^* P_G^f$ we don’t have to think of G as a well-established kind. A third counterargument is the fact that although the stereotype is that Hindus don’t eat meat, a generic like ‘Hindus eat meat’ can be true in certain contexts, e.g. as a denial of the claim that no Hindu eats meat. We have seen above how to deal with problem, however. We conclude that the standard arguments against an analysis of generics in terms of stereotypicality are not valid on our implementation of the latter notion.

Sterken (2015) has recently argued that generics are more context dependent than is generally assumed: not only the domain of quantification is context dependent, but also the required force of quantification. We see her proposal as a description of what should be done. Seen from this perspective, one can see our analysis as a proposal of how this idea should be implemented. Indeed, on our analysis the required force of ‘quantification’ depends on context as well. How high $P(f/G)$ must be in order for the generic ‘ G s are f ’ to be true depends on what $Alt(G)$, and thus $P(f/\bigcup AltG)$ is, what $Value(f)$ is and what $Alt(f)$ is. Sterken (2015) also argues that, in contrast to most people’s intuition, striking generics like (10-b) and (21) are actually false. Such a move certainly makes it much easier to provide an adequate analysis of the truth conditions of generics. Still, one presumably has to make use of something like *Value* anyway to explain the common intuition that such sentences are true, either in semantics or in pragmatics. In this paper we follow our own intuition and the standard assumption in the literature that sentences like (10-b) and (21) are true. See Sect. 5 of this paper, however, for a pragmatic explanation Sterken might use of how high *Value* might give rise to high subjective probability. Given the similarity of generics and habituals and Carlson’s (1977) observation that we need far fewer instances of John murdering children to make the habitual (24-a) acceptable than instances of John walking to work to make (24-b) true,

- (24) a. John murders children.
b. John walks to work.

we advice strongly against such a move, though.

Tessler and Goodman (2019) have very recently proposed another analysis of generics where the value that $P(f/G)$ has to be in order for the generic to be acceptable is context dependent, in particular on prior probabilities. We only discovered this proposal after writing the first versions of this paper, and the analysis is based on similar ideas as we have. For instance, they use (in the latest version) $\frac{P(f/G)}{P(f)}$ as a relevant measure, a notion we have seen was used by others to measure stereotypicality. We don’t

want to compare the analyses in detail here,³⁰ but let us point to at least two salient differences. First, in contrast to us they claim that generics should be accounted for in terms of *subjective* probabilities, and that generics are not simply true or false. Instead, we agree with most semanticists that generics do have truth conditions, for instance to account for the fact that generics can be embedded in, for instance, antecedents of conditionals (cf. Pelletier and Asher 1997). Moreover, we agree with Krifka et al. (1995) that ‘Snakes are slimy’ is false even if most people believe it is true. Second, Tessler and Goodman (2019) don’t make use of *Value* to account for striking generics. They seek to account for these by suitable choices of alternatives. Although we have to admit that the use of *Value* only adds an extra free parameter to the analysis, we feel that the use of this parameter is well justified by (i) other work in formal semantics that make use of utility as well (e.g. van Rooij 2003; Malamud 2012), and (ii) the fact that fear and intensity is known to affect the learning and representation of categories (as reviewed in Sect. 3 of this paper).

Finally, it has been suggested to us by several people that sentences like (10-a) ‘Mosquitos carry the West Nile Virus’ are generally taken to be true not because people think $P(f/G)$ is high, but rather because they take $P(G/f)$ to be high; people just confuse the two kinds of conditional probabilities. There is indeed empirical evidence that people sometimes confuse the two types of conditional probabilities. We doubt, however, that this can be the reason that we consider other ‘striking’ generics like (9) ‘Wolves kill men’ to be true. So although this suggested strategy might be able to explain away some ‘striking’ generics, it won’t explain away all of them, or so we think.

5 Strong pragmatics: from biases to probabilities

On the basis of experimental evidence, Cimpian et al. (2010) concluded that to *accept* a generic about a group we are familiar with, relatively weak conditions have to be fulfilled. At the same time, Cimpian et al. (2010) have observed that hearers treat generics as being inferentially much more powerful. They are *treated inferentially* in a much stronger way: (almost) all *G*s are *f*.³¹ This holds especially if this generic is about a, for the hearer, relatively unknown group. What could explain this strong interpretation?

Our **proposal** is that this is due to the fact that people generally **confuse** representativeness (or stereotypicality) with probability (or prototypicality), especially if they are less familiar with *G*s. Thus, for those cases where $Alt(G) \neq \emptyset$, and thus $\Delta^* P_G^f$ does not reduce to $P(f/G)$, we propose that, nevertheless, people confuse high $\Delta^* P_G^f$ with high $P(f/G)$. This idea might seem ad hoc, but it is in fact at the heart of the

³⁰ The use of $\frac{P(f/G)}{P(f)}$ instead of our $\Delta^* P_G^f$ will have some empirical consequences, because the notions are both quantitatively and ordinally distinct. At this point we don’t know how important these empirical consequences are.

³¹ This distinction between conditions for acceptance and inferential use (or perhaps understanding versus willingness to produce) is arguably not limited to generics. A similar distinction has been proposed for the analysis of, among others, sentences involving (negated) vague adjectives, e.g., by Krifka (2007) and Cobreros et al. (2012).

whole Biases and Heuristics program of Tversky and Kahneman (1974). We will argue that people confuse high ΔP_G^f or high $\Delta^* P_G^f$ with high conditional probability due to their representativeness and causality heuristics (or close variants thereof), and that they give higher probabilities to features with high *Value* due to their availability heuristics. As a result, high $\nabla P_G^f = \frac{\Delta P_G^f}{1-P(f/-G)} \times \text{Value}(f)$ is confused with high $P(f/G)$. This, we propose, gives rise to the strong inferential interpretation.

The heuristics and biases program started with Tversky and Kahneman showing that our intuitions involving probability judgements are not in accordance with the norms given by Bayesian probability theory. Kahneman and Tversky's (1972) most famous example is their conjunction fallacy, which shows that in some situations people assign greater (conditional) probability to a conjunction than to one of its conjuncts, i.e., $P(B \wedge F/L) > P(B/L)$, although this is impossible according to the normative Bayesian theory. For example, a woman (Linda) with liberal political views was judged by most participants to be more likely a feminist bank teller than a bank teller.

Tversky and Kahneman did not only point out some empirical problems for the normative theory of probability. They also proposed an alternative hypothesis of how we actually make probability judgements. According to their Biases and Heuristics program (Tversky and Kahneman 1974), to reach a probability judgement, we often do not reason according to Bayesian probability theory, but use simplifying or shortcut heuristics. These heuristics are mostly approximately correct, but also give rise to systematic biases in certain contexts.

At the heart of the Heuristics and Biases approach is the attribute substitution heuristic, and biases like representativeness, availability, and causality. The heuristic people make use of is that when faced with a hard question about a particular quantity, type of situation or attribute, people have a tendency to answer a different but easier question about quantities, types of situations or attributes that are still *representative* for the ones in the original question, but are also more readily accessible (cf. Kahneman and Frederick 2002). Kahneman and Tversky (1972) argue that the conjunction fallacy arises because individuals apply the representativeness heuristics as a cue to subjective probability, $B \wedge F$ (being a bank teller *and* a feminist) given L is considered more likely than B (being a bank teller) given L , because $B \wedge F$ is considered more typical, or representative, for L inda.

Kahneman and Tversky (1972) thought of representativeness in terms of a primitive notion of *similarity*, but this has been criticized as being too vague to explain our intuitive probability judgements (cf. Gigerenzer 1996). Tenenbaum and Griffiths (2001) and others convincingly argued that for conceptual (rationalistic) and empirical reasons, representativeness can better be thought of in terms of a notion of evidential support, or associative strength as measured by something like ΔP or $\Delta^* P$ or $(\log) \frac{P(f/G)}{P(f/-G)}$.³² Indeed, Gluck and Bower (1988), Shanks (1990) and Lagnado and Shanks (2002) explicitly claim that when people are asked a question about proba-

³² Although these measures differ from one another, they have a lot in common. Most importantly, these values go up with increasing difference between $P(f/G)$ and $P(f/-G)$. All three notions, and many similar ones as well, have also been proposed in philosophy of science to measure strength of confirmation.

bility, they readily substitute this with the closely related question about evidential support, or strength of association. Although they claim that this makes a lot of sense (because to be attuned to evidential strength, or support, is in general more important than to be attuned to absolute probability values), it sometimes gives rise to a very incorrect response.³³ To give a telling example from Newel et al. (2007), suppose that a football team is as likely to win as to lose when Johan plays, but that the team much more likely loses when Johan is not playing. In that case, although $P(\text{win}/\text{Johan plays}) = P(\text{-win}/\text{Johan plays})$, still people typically will believe that the team will win if Johan is playing. Indeed, Lagnado and Shanks (2002), Crupi et al. (2008) and others show that with the help of any of the notions of evidential support, the conjunction fallacy can be explained³⁴: although $P(B \wedge F/L) \leq P(B/L)$, it might well be that $P(B \wedge F/L) - P(B \wedge F/\neg L) > P(B/L) - P(B/\neg L)$ and thus $\Delta^{(*)} P_L^{B \wedge F} > \Delta^{(*)} P_L^B$. In a similar way one can explain other ways people deviate from the normative Bayesian theory, such as the fact that people tend to neglect base rates (see e.g., Newel et al. 2007). More in general, the experimental evidence that people are more accurate and consistent in making impact judgments than probability judgments suggest that ‘stereotypes could be generated and maintained by a prevalence of impact over posterior probability’ and ‘people appear rational because they rely more on detecting relations of impact than on computing values of posterior probability’ (Tentori et al. 2016, p. 770).

The problem we wanted to account for in this section is to explain why people generally treat generics of the form ‘Gs are f ’ inferentially as holding that $P(f/G)$ is high. Our analysis merely predicts that the sentence is true iff ∇P_G^f is high, which means that $\Delta^* P_G^f \times \text{Value}(f)$ is high. On the above idea (known as the ‘associative theory of probability judgments’) of Gluck and Bower (1988) and others that we readily substitute, or confuse, probability for evidential support, the gap between the two can easily be bridged in case all features have the same value. Recall that if the value of the features is irrelevant, high ∇P_G^f reduces to a high relative difference, $\Delta^* P_G^f$. By the associative theory of probability judgments, or the representativeness bias, however,

³³ More formally, whereas contingency and relative difference are symmetric-like notions, conditional probability is not. $P(f/G)$ can increase without $P(G/f)$ doing so. On the other hand, we have seen above that ΔP_G^f behaves monotone increasingly with respect to $P(f/G) - P(f)$. To show that contingency behaves symmetric-like, it suffices to show that $P(f/G) - P(f)$ behaves monotone increasingly with respect to $P(G/f) - P(G)$:

$$\begin{aligned} P(f/G) - P(f) &= \frac{P(G/f) \times P(f)}{P(G)} - \frac{P(G) \times P(f)}{P(G)} \\ &= \alpha [P(G/f) \times P(f)] - \alpha [P(G) \times P(f)], \text{ with } \alpha = \frac{1}{\gamma}, 0 \leq \gamma \leq 1 \\ &= \alpha [P(G/f) \times P(f) - P(G) \times P(f)] \\ &= \alpha [P(G/f) \times \beta - P(G) \times \beta], \text{ with } 0 \leq \beta \leq 1 \\ &= \alpha \beta [P(G/f) - P(G)]. \end{aligned}$$

³⁴ To be clear, there is no shortage of proposals of how to account for the conjunction fallacy. Some have proposed to make use of likelihood, of source reliability, or of a nonmonotonic logic analysis of normality, instead of evidential support. Others have pointed to the potential influence of a non-Boolean reading of conjunction, of conversational implicatures, of alternatives being considered, or of cognitive ability. Many empirical studies have been done to control for these influences, but Tversky and Kahneman’s findings proved to be rather robust. If the participants were confronted explicitly with frequencies instead of probabilities, however, the conjunction error reduced sharply.

this is confused with high $P(f/G)$, which explains the common intuition under these circumstances.

What if the values of the features might be important, i.e., if the relevant features are ‘horrific’ or ‘appalling’? According to Tversky and Kahneman’s (1974) *availability heuristics*, people assess the probability of an event by the ease with which instances or occurrences can be brought to mind. Usually this heuristic works quite well; all things being equal, common events are easier to remember or imagine than uncommon events. Unfortunately, sometimes the general rule of thumb doesn’t do its job and leads to systematic biases. Some events are more available than others not because they tend to occur frequently or with high probability, but simply because they are inherently easier to think about. Emotionally-charged events, or horrific or appalling features, can be brought to mind easily. Thus, by the availability heuristics, $P(f/G)$ is considered higher than it actually is, perhaps close to 1, if $Value(f)$ is high.

We have argued that hearers that are unfamiliar with group G treat a generic of the form ‘ G s are f ’ inferentially in a much stronger way than one would expect according to the semantic analysis we defended in the previous section, because they confuse representative features, features f where $\Delta^*P_G^f \times Value(f)$ is high, with probable features, features f where $P^*(f/G)$ is high. In the last formula, P^* models subjective belief rather than objective frequencies or propensities. Obviously, if $P^*(f/G)$ is high, we have explained why the generic ‘ G s are f ’ is interpreted as meaning that almost all G s are f .

Actually, there is another reason why a high representativeness between G and f , measured by a high ΔP_G^f or high $\Delta^*P_G^f$, increases the felt conditional probability $P(f/G)$, i.e., $P^*(f/G)$. Although ΔP_G^f and $\Delta^*P_G^f$ only measure the correlation between G and f , psychologists and philosophers of science use ΔP_G^f or $\Delta^*P_G^f$ as an indication for the causal strength of G for f (e.g., Skyrms 1980; Shanks 1995; Cheng 1997). Indeed, Tversky and Kahneman (1980) show that if we see a correlation, we tend to interpret it in the preferred (strongest) way: as *causal*. Moreover, they show that if an event is causally explained, it is seen as *more likely* than it actually is. Combining these ideas leads to the causality bias: high representativeness of a feature for G as we measured it, leads to higher subjective probability given G than is justified (meaning that $P^*(f/G)$ is higher than $P(f/G)$), because we believe that there exists a causal relation between G and f .

This last reason why high representability leads to high conditional probability is closely related with Barth’s (1971), Leslie’s (2013) and Haslanger’s (2010, 2014) proposals why we (and children in particular) *treat* generics *inferentially* typically in such a strong way. Haslanger (2014) argues that if we use a generic like ‘Women are more nurturing than men’, we implicate that there is something about what it is to be a woman and about what it is to be a man that explains their supposed differential capacities to nurture. From this she concludes that the utterance of a generic of the form ‘ G s are f ’ will normally add to the common ground of the conversation a claim about f ’s naturalness for the group, or kind, G . Why? Because we (perhaps erroneously) believe that robust regularities are not accidental: they are due to the natures of things. This, again, is exactly in line with the view of Keil (1989) and others, that people tend to have essentialist beliefs about a substantial number of categories. Haslanger (2014)

and Leslie (2013) argue that this is why it is dangerous to claim ‘Muslims are terrorists’ but not (10-b) ‘Ticks carry the Lyme disease’: while for the latter essentialists belief might be true, this is certainly not the case for the former.³⁵

Although we agree with Barth (1971), Leslie (2013) and Haslanger (2014) that essentialist beliefs play a pragmatically significant role in why we interpret generic statements in such a strong way (cf. van Rooij and Schulz, to appear), we don’t think this is the whole reason: it is only one of the biases singled out by Tversky and Kahneman (1974) that are crucial.

6 Conclusion and outlook

In this paper we have based our analysis of generic sentences primarily on an intuition that some authors over the years have claimed would be natural for at least some examples (e.g. Krifka et al. 1995): a generic of the form ‘Gs are f ’ is true iff f is a typical, distinguishing, feature for Gs. Many analyses of generics have been proposed over the years, and none has come out as the clear winner. This is partly due, we suspect, to the vagueness and context-dependence of what is meant by a generic. We have little doubt that our proposal won’t meet universal acceptance either. Still, we hope that this paper at least shows that an analysis in terms of typicality can be pushed much further than is generally assumed. We also argued that such a semantic analysis is naturally extended by pragmatic strengthening, making use of insights from Tversky and Kahneman’s Heuristics and Biases approach. This popular approach within social and cognitive psychology [as measured by the selling rates of Kahneman (2011)], has, to the best of our knowledge, never been used so far in pragmatics. We think this is a shame, and we hope this paper will help to change things accordingly.

In this paper we make use of typicality based on contingency, ΔP_G^f , or relative difference, $\Delta^* P_G^f$, to explain why people *accept* generics based on weak evidence. We also use the Heuristics and Biases approach to explain why people *treat* generic statements *inferentially* in a much stronger way. Both ideas are closely related: a high contingency between two variables C and O means that C is *predictive* of O , which suggests a high conditional probability of O given C (and this suggestion might be incorrect in case the base rate of C , $P(C)$, is relatively low). Although the Heuristics and Biases approach and the idea that people readily confuse conditional probability with evidential impact have a solid experimental base, direct experimental evidence of our use of representativeness for the proposed semantic and pragmatic analyses of generics is still missing. But it should be relatively easy to test these proposals experimentally, certainly if representativeness is reduced to relative difference. If there is a conflict between the relative difference (high) and the conditional probability/frequencies (relatively low) between two variables (say G and f), would subjects accept the generic ‘Gs are f ’ if $Alt(G) \neq \emptyset$? And if yes, would these subjects assign (much) higher conditional probabilities to f given G than the actual frequencies allow for? Although we are not yet ready to state definite conclusions, the early results of

³⁵ Moreover, $P(\text{terrorist}/\text{Muslims})$ is many orders of magnitudes lower than $P(\text{carry Lyme disease}/\text{Ticks})$.

some online experiments where we tested these hypotheses do, indeed, suggest that the answers to those questions are positive [see also the experiments on stereotypes as reported in Bordalo et al. (2016)]. In this paper we also suggested that subjects often accept and interpret generics not based on actual frequencies, or propensities, but on a distorted picture of it provided by the media. This suggestion was investigated by our student van Harmelen (2017) who tested (on a relatively small corpus) whether there is a relation between acceptance of generic sentences of the form ‘ G s are f ’ and the contingency between the two variables as observed in newspapers, making use of techniques from distributional semantics. His findings are encouraging.

In this paper we have linked generics with typicality, and in particular stereotypicality. Thinking of *stereotypes* in terms of representativeness makes them intrinsically useful, because they allow us to make generalizations. But, of course, stereotypes have a negative connotation: they misrepresent groups for which they are typical. We have explained how that can be: representativeness, or typicality, normally looks for distinguishing features, and as such, exaggerate differences. But we haven’t explained many other properties of stereotypes (cf. Schneider 2004), such as why stereotypes—and thus the acceptance of generic sentences—are so hard to change. We think we can explain this as well, but for that we would need stronger representations of beliefs. Until now we have represented beliefs in terms of probabilities, but we haven’t been able to model their difference in stability. Perhaps this could be accounted for in terms of a connectionist mechanism to learn ∇P , making use of the delta rule which is sensitive to the sample size. Alternatively, we could make use of Skyrms’ (1980) notion of ‘resilience’ to capture the notion of stability, although perhaps causal frameworks (e.g. Pearl 1988, 2000) could provide a more insightful model. In future work we would like to investigate how best to proceed.

Our analysis of generics in this paper is very Humean, built on correlations and frequencies and the way we learn from those. Many linguists and philosophers feel that there must be something more: something underlying these actual frequencies, like *essences* or objective *kinds* which have *causal powers* and *dispositions*. In future work we hope to show that such a causal perspective is closely related to the associative analysis of generics proposed in this paper, making use of insights due to Cheng (1997) and Pearl (2000) and others. We believe something similar holds for *conditionals*. It seems a natural thought to analyze generics as hidden conditionals, or at least for the two to be handled in similar ways. Indeed, similarly to generics, for the conditional ‘If A , then C ’ to be appropriate it is normally claimed that it is high $P(C/A)$ that counts (Adams 1965). But high $P(C/A)$ can hardly be enough for all conditionals, in particular not for the so-called ‘relevance’ conditionals. Moreover, although the condition that $P(C/A)$ is high is fulfilled for the conditional ‘If A , then $0 \neq 1$ ’, it is still a weird, if not inappropriate, conditional. Some connection is desired, normally a *causal* one. According to Shanks (1995) and Cheng (1997) and others, positive ΔP_A^C is a necessary condition for A being causally relevant to C . Thus, it is only natural to demand that ΔP_A^C , and thus ∇P_A^C , is positive. In further work we will propose that for

the appropriateness of many conditionals, what counts is high ∇P_A^C , perhaps derived from a causal analysis of conditionals (cf. van Rooij and Schulz 2019).

Funding The funding was provided by FP7 Ideas: European Research Council (Grant No. Marie Curie Actions—Initial Training Networks (ITN)).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Adams, E. W. (1965). A logic of conditionals. *Inquiry*, 8, 166–197.
- Ameel, E., & Storms, G. (2006). From prototypes to caricatures: Geometrical models for concept typicality. *Journal of Memory and Language*, 55, 402–421.
- Asher, N., & Morreau, M. (1995). What some generic sentences mean. In G. Carlson & F. Pelletier (Eds.), *The generic book* (pp. 300–339). Chicago: University of Chicago Press.
- Barsalou, L. W. (1985). Ideals, central tendency and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 11, 629–654.
- Barth, E. (1971). *De Logica van de Lidwoorden in de Traditionele Filosofie*. Leiden: Universitaire Pers Leiden.
- Bernardo, J. M. (1979). Expected information as expected utility. *Annals of Statistics*, 7, 686–690.
- Bordalo, P., Coffman, K., Gennaioli, N., & Shleifer, A. (2016). Stereotypes. *Quarterly Journal of Economics*, 131, 1753–1794.
- Burnett, R., Medin, C., Ross, N., & Blok, S. (2005). Ideal is typical. *Canadian Journal of Experimental Psychology*, 59, 3–10.
- Cahill, L., Babinsky, R., Markowitsch, H., & McGaugh, J. (1995). The amygdala and emotional memory. *Nature*, 377, 295–296.
- Carlson, G. (1977). *Reference to kinds in English*. Ph.D. dissertation, University of Massachusetts, Amherst.
- Chatlosh, D., Neunaber, D., & Wasserman, E. (1985). Response-outcome contingency: Behavioral and judgmental effects of appetitive and aversive outcomes with college students. *Learning and Motivation*, 16, 1–34.
- Cheng, P., & Holyoak, K. (1995). Adaptive systems as intuitive statisticians: Causality, contingency, and prediction. In J. Meyer & H. Roitblat (Eds.), *Comparative approaches to cognition*. Cambridge, MA: MIT Press.
- Cheng, P. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367–405.
- Cimpian, A., Brandone, A., & Gelman, S. (2010). Generic statements require little evidence for acceptance but have powerful implications. *Cognitive Science*, 34, 1452–1482.
- Cobrerros, P., Egré, P., Ripley, D., & van Rooij, R. (2012). Tolerant, classical, strict. *Journal of Philosophical Logic*, 2012(41), 347–385.
- Cohen, A. (1999). *Think generic! The meaning and use of generic sentences*. Stanford: CSLI Publications.
- Cohen, A. (2001). On the generic use of indefinite singulars. *Journal of Semantics*, 18, 183–209.
- Cohen, A. (2004a). Existential generics. *Linguistics and Philosophy*, 27, 137–168.
- Cohen, A. (2004b). Generics and mental representations. *Linguistics and Philosophy*, 27, 529–556.
- Crupi, V., Fitelson, B., & Tentori, K. (2008). Probability, confirmation, and the conjunction fallacy. *Thinking and Reasoning*, 14, 182–199.
- Declerck, R. (1996). The manifold interpretation of generic sentences. *Lingua*, 68, 149–188.
- Dunbar, R. (1996). *Grooming, gossip and the evolution of language*. London: Faber and Faber.
- Eckardt, R. (1999). Normal objects, normal worlds, and the meaning of generic sentences. *Journal of Semantics*, 16, 237–278.
- Fernando, T., & Kamp, H. (1996). Expecting many. In T. Galloway & J. Spence (Eds.), *Proceedings of the sixth conference on semantic and linguistic theory* (pp. 53–68). Ithaca, NY: CLC Publications.

- Geurts, B. (1985). Generics. *Journal of Semantics*, 4, 247–255.
- Gigerenzer, G. (1996). On narrow norms and vague heuristics. *Psychological Review*, 103, 592–596.
- Good, I. (1950). *Probability and the Weighing of Evidence*. London: Griffin.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, 117, 227–247.
- Haslanger, S. (2010). Ideology, generics, and common ground. In C. Witt (Ed.), *Feminist metaphysics: Essays on the ontology of sex, gender and the self* (pp. 179–207). Dordrecht: Springer.
- Haslanger, S. (2014). The normal, the natural and the good: Generics and ideology. *Politica & Societa*, 3, 365–392.
- Hoeltje, M. (2017). Generics and ways of being normal. *Linguistics and Philosophy*, 40, 101–118.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3, 430–454.
- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases* (pp. 49–81). New York: Cambridge University Press.
- Kahneman, D. (2011). *Thinking, fast and slow*. Straus and Giroux, New York: Farrar.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: Bradford Books/MIT Press.
- Klein, E. (1980). The semantics of positive and comparative adjectives. *Linguistics and Philosophy*, 4, 1–45.
- Krifka, M., Pelletier, F., Carlson, G., ter Meulen, A., Chierchia, G., & Link, G. (1995). Genericity: An introduction. In G. Carlson & F. Pelletier (Eds.), *The generic book* (pp. 1–124). Chicago: University of Chicago Press.
- Krifka, M. (2007). Negated anonyms: Creating and filling the gap. In U. Sauerland & P. Stateva (Eds.), *Presupposition and implicature in compositional semantics* (pp. 163–177). New York: Palgrave Macmillan.
- Krifka, M. (2013). Definitional generics. In A. Mari, C. Beyssade, & F. Del Prete (Eds.), *Genericity* (pp. 372–389). Oxford: Oxford University Press.
- Kripke, S. (1972/80). Naming and necessity. In: D. Davidson & G. Harman (Eds.), *Semantics of natural language* (pp. 253–355, 763–769). Dordrecht: Reidel.
- Lagnado, D. A., & Shanks, D. (2002). Probability judgment in hierarchical learning: A conflict between predictiveness and coherence. *Cognition*, 83, 81–112.
- Lawler, J. (1973). *Studies in English generics*, University of Michigan Papers in Linguistics.
- Leslie, S. J. (2008). Generics: Cognition and acquisition. *The Philosophical Review*, 117, 1–47.
- Leslie, S. J. (2013). Essence and natural kinds: When science meets preschooler intuition. In T. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 4, pp. 108–166). Oxford: Oxford University Press.
- Leslie, S. J., Khemlani, S., & Glucksberg, S. (2011). Do all ducks lay eggs? The generic overgeneralization effect. *Journal of Memory and Language*, 65, 15–31.
- Lynch, E. B., Coley, J. D., & Medin, D. L. (2000). Tall is typical: Central tendency, ideal dimensions, and graded category structure among tree experts and novices. *Memory & Cognition*, 28, 41–50.
- Malamud, S. (2012). The meaning of plural definites: A decision-theoretic approach. *Semantics and Pragmatics*, 5, 1–48.
- McCarthy, J. (1986). Application of circumscription to formalize common-sense knowledge. *Artificial Intelligence*, 28, 89–116.
- McCauley, C., Stitt, C. L., & Segal, M. (1980). Stereotyping: From prejudice to prediction. *Psychological Bulletin*, 87, 195–208.
- Newel, B., Lagnado, D., & Shanks, D. (2007). *Straight Choices. The psychology of decision making*. Hove and New York: Psychology Press.
- Nickel, B. (2009). Generics and the ways of normality. *Linguistics and Philosophy*, 31, 629–648.
- Nickel, B. (2012). Dutchmen are good sailors: Generics and gradability. In A. Mari, C. Beyssade, & F. Del Prete (Eds.), *Genericity* (pp. 390–405). Oxford: Oxford University Press.
- Palmeri, T., & Nosofsky, R. (2001). Central tendencies, extreme points, and prototype enhancement effect in ill-defined perceptual categorization. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 54, 197–235.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. San Mateo, CA: Morgan Kaufman Publishers, Inc.
- Pearl, J. (2000). *Causality: Models, Reasoning and Inference*. Cambridge: Cambridge University Press.

- Pelletier, F., & Asher, N. (1997). Generics and defaults. In J. van Benthem & A. ter Meulen (Eds.), *Handbook of logic and language* (pp. 1125–1177). Amsterdam: North Holland.
- Pierce, W. D., & Cheney, C. D. (2004). *Behavior analysis and learning* (3rd ed.). Mahwah, NJ: Lawrence Erlbaum Associates Publishers.
- Reiter, R. (1987). Nonmonotonic reasoning. In J. F. Traub, N. J. Nilsson, & B. J. Grosz (Eds.), *Annual review of computer science* (pp. 147–186). Palo Alto: Annual Reviews Inc.
- Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, *66*, 15.
- Rescorla, R., & Wagner, A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. Black & W. Prokasy (Eds.), *Classical conditioning II* (pp. 64–99). New York: Appleton-Century-Crofts.
- Rosch, E. (1973). Natural categories. *Cognitive Psychology*, *4*, 328–350.
- Schneider, D. (2004). *The psychology of stereotyping*. New York, NY: The Guilford Press.
- Schubert, L., & Pelletier, F. (1987). Problems in the representation of the logical form of generics, plurals and mass nouns. In E. LePore (Ed.), *New direction in semantics* (pp. 357–371). London: Academic Press.
- Seeley, T. (1985). *Honeybee ecology: a study of adaptation in social life*. Princeton: Princeton University Press.
- Shanks, D. (1990). Connectionism and the learning of probabilistic concepts. *Quarterly Journal of Experimental Psychology*, *42A*, 209–237.
- Shanks, D. (1995). *The psychology of associative learning*. Cambridge: Cambridge University Press.
- Shep, M. C. (1958). Shall we count the living or the dead? *New England Journal of Medicine*, *259*, 1210–1214.
- Skyrms, D. B. (1980). *Causal Necessity. A pragmatic investigation of the necessity of laws*. New Haven and London: Yale University Press.
- Slovic, P., Finucane, M., Peters, E., & MacGregor, D. G. (2004). Risk as analysis and risk as feelings: Some thoughts about affect, reason, risk, and rationality. *Risk Analysis*, *24*, 1–12.
- Stalnaker, R. C. (1978). Assertion. In P. Cole (Ed.), *Syntax and Semantics: Pragmatics* (Vol. 9, pp. 315–332). New York: Academic Press.
- Sterken, R. K. (2015). Generics in context. *Philosopher's Imprint*, *15*, 1–30.
- Tenenbaum, J., & Griffiths, T. (2001). The rational basis of representativeness. In J. Moore & K. Stenning (Eds.), *Proceedings of the 23th annual conference of the cognitive science society* (pp. 1036–1041).
- Tentori, N., Chater, N., & Crupi, V. (2016). Judging the probability of hypotheses versus the impact of evidence. *Cognitive Science*, *40*, 758–778.
- Tessler, M., & Goodman, N. (2019). The language of generalization. *Psychological Review*, *126*(3), 395–436.
- Thomas, L., & LaBar, K. (2008). Fear relevancy, strategy use, and probabilistic learning of cue-outcome associations. *Learning and Memory*, *15*, 777–784.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124–1131.
- Tversky, A., & Kahneman, D. (1980). Causal schemas in judgments under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology* (pp. 49–72). Hillsdale, NJ: Erlbaum.
- van Harmelen, M. (2017). *On the acceptance of generic sentences*. Bachelor thesis, University of Amsterdam.
- van Rooij, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy*, *26*, 727–763.
- van Rooij, R. (2004). Utility, informativity and protocols. *Journal of Philosophical Logic*, *33*, 389–419.
- van Rooij, R., & Schulz, K. (2019). Conditionals, causality and conditional probability. *Journal of Logic, Language and Information*, *28*, 55–71.
- van Rooij, R., & Schulz, K. (to appear). A causal power theory of generics. *Topoi: An International Review of Philosophy*.
- Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, *25*, 221–261.

Affiliations

Robert van Rooij¹  · Katrin Schulz¹

Katrin Schulz
K.Schulz@uva.nl

¹ Institute for Logic, Language and Computation (ILLC), University of Amsterdam, Amsterdam, The Netherlands