

Reconceptualising the Psychological Theory of Generics

Tom Ralston

Final version to appear in *Philosophical Studies*

Abstract

Generics have historically proven difficult to analyse using the tools of formal semantics. In this paper, I argue that an influential theory of the meaning of generics due to Sarah-Jane Leslie, the Psychological Theory of Generics, is best interpreted not as a theory of their meaning, but as a theory of the psychological heuristics that we use to judge whether or not generics are true. I argue that Leslie's methodology is not well-suited to producing a theory of the meaning of generics, since it takes speakers' judgments at face value and ignores the non-semantic factors that might affect these judgments. Leslie's theory therefore overfits the data of our linguistic intuitions. I present a reconceptualised version of the Psychological Theory of Generics as a theory of how heuristics affect our judgements of the truth values of generics and discuss the application of this reconceptualised theory to some of the puzzles posed by generics, including their apparent content-sensitivity, their inferential asymmetry and their association with stereotyping and prejudice.

Keywords: Generics, psychological theory of generics, heuristics and biases, semantics

1 Introduction

One puzzle about the semantics of generics is that our intuitive judgments of their truth values do not seem to track facts about the prevalence of a property among members of a category, which has made providing an account of the semantics of generics in set-theoretic terms extremely challenging. The Psychological Theory of Generics (PTG) was developed by Sarah-Jane Leslie (2007; 2008) to account for this fact. According to PTG, generics express a cognitively default mode of

generalisation, which is why our intuitions regarding the truth values of generics are so varied.

Picking up on a line of argument raised by van Rooij and K. Schulz (2020), I will argue that PTG is best interpreted not as a theory of the meaning of generics, but as a theory of the psychological heuristics that we use to judge whether or not generics are true.

The paper begins with a description of PTG and a discussion of the risks associated with interpreting PTG as a theory of the meaning of generics. I then argue that Leslie’s methodology is not well-suited to producing a theory of the meaning of generics, on the basis that it takes speakers’ judgments at face value and ignores the non-semantic factors that might affect these judgments. Finally, I present a reconceptualised version of PTG as a theory of how heuristics affect our judgements of the truth values of generics and discuss the application of this reconceptualised theory to some of the puzzles posed by generics, including their apparent content-sensitivity, their inferential asymmetry and their association with stereotyping and prejudice.

1.1 The Psychological Theory of Generics

I will focus, as Leslie does, on generic readings of English bare plural sentences (‘generics’), which are sentences like the following:

- (1) Swiss men perform military service.
- (2) Dogs have four legs.

The generic reading of a bare plural sentence, ‘Ks are F’, does not entail that As are F when all Ks are As (Leslie, 2008: 5 n. 2). This distinguishes generic readings from existential readings of bare plurals. For example, ‘Swiss men are standing over there’ only has an existential reading, since it only has a reading which entails that men are standing over there.

Generics have historically proven difficult to analyse using the tools of formal semantics.¹ Because of this, Leslie takes a heterodox approach to the study of the meaning of generics. She is, she explains, not interested in the traditional semantic question of how our words relate to the world, but to the psycholinguistic question of how “human beings actually understand their natural languages” (2008: 20). She therefore forgoes the standard tools of formal semantics in favour of her psychological approach. The central idea behind PTG is that humans are hard-wired to generalise in a certain way, and this mode of generalisation is not particularly

¹There are several surveys of approaches to the semantics of generics: Krifka et al. (1995), Mari et al. (2013), Sterken (2017), Nickel (2017), and Leslie and Lerner (2022).

systematic. Generics give voice to this mode of generalisation, which explains the perplexing features of their semantics.

Eschewing a metaphysical account of genericity, Leslie builds PTG around the idea that generics are the expression of the cognitive system’s “default” mode of generalisation. This mode of generalisation is not linguistic but predates the acquisition of language in cognitive development (Leslie, 2007: 381). This explains the “near-paradoxical nature of the acquisition of generics” (Leslie, 2008: 16): children develop the ability to use and understand generics from a very early age, before they learn to use explicit quantifiers (ibid.: 19). Leslie’s hypothesis neatly explains why this should be the case, since children have an innate ability to generalise in this ‘default’ manner, the linguistic expression of which is a generic.

This mode of generalisation is closely associated with our “System 1” reasoning, which is fast and insensitive to factors such as probability and quantity, drawing conclusions on the basis of heuristics and biases (Leslie, 2007: 396; cf. Kahneman, 2003: 688-689). System 2, by contrast, is slower but more precise and is, according to Leslie, more closely associated with quantified generalisations whose semantics can be stated in set-theoretic terms.

PTG sorts generics into various categories which have different, though related, truth-conditions. The different categories derive from the content of the generic itself. PTG is therefore attributes highly *content-sensitive* truth-conditions to generics. One class of troublesome examples which are an important motivation for Leslie’s theory, and have formed the basis of much of her subsequent work on generics (e.g. Leslie, 2017) are referred to as ‘striking property’ generics. These are generics which attribute a striking or dangerous property to a kind:

- (3) Mosquitoes carry the West Nile Virus.
- (4) Sharks attack bathers (Leslie, 2008: 14).

Leslie maintains that these generics are true, even though the vast majority of mosquitoes do not carry the West Nile Virus and the vast majority of sharks never attack bathers. This view receives some empirical support from a truth value judgment task and an implied prevalence task in Cimpian, Brandone, et al. (2010: Experiment 1).

Leslie argues that there are a variety of ways in which generics can be true. She offers the following description of what she calls the “worldly truth specifications” of generics:

We can describe the circumstances under which a generic of the form ‘Ks are F’ is true as follows:

The counterinstances are negative, and:

If F lies along a characteristic dimension for the Ks, then some Ks are F, unless K is an artifact or social kind, in which case F is the function or purpose of the kind K;

If F is striking, then some Ks are F and the others are disposed to be F;

Otherwise, almost all Ks are F (Leslie, 2008: 43).

It is worth noting immediately that Leslie does not present these conditions as part of the compositional semantics of generics. They provide a metaphysical account of “how the world must be for the sentence to be true” (ibid.: 43). She regards it as unlikely that we will be able to give an informative semantics for generics which does not use the generic operator Gen, the object-language expression of which is not phonologically realised, in the metalanguage, just as we state the truth-conditions for a sentence like ‘Bob is red’ by using the predicate ‘is red’ in the metalanguage: ‘Bob is red’ is true iff Bob is red (Leslie, 2007: 387-388). So she states the (semantic) truth-conditions of a generic such as (4) in the following way:

(5) ‘Sharks attack bathers’ is true iff $\text{Gen}(x) : \text{shark}(x) [\text{attack bathers}(x)]$

The lengthy ‘worldly truth specifications’ above are an attempt to spell out what the world has to be like in order for the right-hand side of this biconditional to be true, a task which, she argues, takes us beyond semantics. Regardless of whether it makes sense to talk, as Leslie does, about a distinction between “worldly truth specifications” and “semantically derived truth conditions”,² it is surprising that the *worldly* truth specifications depend not only on mind-independent features of the world, but on psychological phenomena, such as a property’s being “striking” or a counterinstance’s being “negative”. This focus on the psychological, rather than semantic or metaphysical, suggests that the objective of PTG is not to offer a theory of the meaning of generics, since it focuses on the conditions under which we make certain judgments about generics, not the conditions under which they are true.

I won’t discuss the accuracy of Leslie’s account of the truth-conditions of generics in this paper (see Sterken, 2015a for an extended discussion). Instead, I will argue that Leslie makes an error in taking the judgments of speakers about the truth values of generics at face value, and so presents PTG as a theory of the meaning of

²Throughout, I will use ‘truth-conditions’ to mean ‘the conditions under which a sentence is true’. I leave it to the reader to decide whether the semantic truth-conditions of a sentence can differ from its metaphysical truth-conditions.

generics, when in fact the evidence supports a theory which explains how people make judgments about the truth of generics.

1.2 Theory Confusion

As will become clear, the questions that PTG seeks to answer are quite different from those in the formal semantics literature on generics. This is not always made clear; for example, in Sterken’s overview of contemporary theories of generics (2017), “Psychologically Based Theories”, in which only Leslie’s PTG is included, are presented alongside “Normality Theories” and contextualist approaches, as though they are competing theories attempting to answer the same questions. Nickel does the same in his handbook article on generics (2017).³ But if those theories are attempting to provide a compositional semantics on the basis of our intuitive judgments about the truth values of generics, whereas PTG gives an account of when and why people are disposed to judge generics to be true, they are not in competition.⁴ In fact, PTG may well be compatible with a range of semantic theories for generics, provided that the semantic theories are compatible with the psychological explanations PTG offers for our truth value judgments about generics. Some authors, especially those who focus on the variability of the meaning of generics in the way that Leslie does, explicitly acknowledge that PTG is not a competitor to their own semantic theory of generics (e.g. Sterken, 2015a: 27; Lee and Nguyen, 2022: 1757). Leslie’s suggestion that generics are associated with an unsystematic default mode of generalisation might explain why the truth-conditions of generics are so variable, without itself offering a theory of their truth-conditions.⁵ Semantic theories of generics need not, however, vindicate the truth value judgments that PTG is based on, as it is quite possible that people are systematically mistaken in their evaluation of generics (Sterken, 2015c: 2508).

There is a danger that, outside of the formal semantics literature, these fine distinctions are not properly understood, and PTG is accepted as a semantic theory on the basis that it handles the data better than any alternatives. There is a burgeoning literature which attempts to reach conclusions about subjects as diverse

³Nickel does, however, acknowledge that “cognitive approaches to generics generally, and Leslie’s work in particular, raise a host of foundational questions about the relationship between semantics and human cognition” (2017: 456).

⁴In her entry on generics in the Stanford Encyclopedia of Philosophy, Leslie draws the kind of distinction that I propose here and separates her psychological account from “semantic analyses of [the operator] *Gen*” (Leslie and Lerner, 2022). She would, I take it, be sympathetic with the views I put forward here, although the presentation of PTG in her early papers muddies the waters somewhat.

⁵I am grateful to an anonymous reviewer for raising this point.

as Kripkenstein’s scepticism about rule-following (Cheng, 2011), epistemic injustice (Rosola and Cella, 2020), ideology (Haslanger, 2012; Langton et al., 2012), decision theory (Schiller, 2023) and scientific communication (DeJesus et al., 2019; Peters, 2021) on the basis of considerations about the meaning of generics. Those who make normative prescriptions about our use of language based explicitly on the semantics given by PTG, such as Langton et al. (2012) and Claveau and Girard (2019), should be clear that the foundations of their arguments are built on a theory of how we evaluate generics, not what they mean. This need not invalidate their arguments; it is possible that all these arguments require is that we are disposed to evaluate generics as true under certain conditions, not that they actually are true under those conditions.⁶

But failing to heed the distinction between how we think about generics and what generics mean risks placing too much weight on the potential effects of language change and in particular the effects of replacing generics with explicitly quantified generalisations (e.g. Langton et al., 2012: 765). If humans are hard-wired to think in the way that Leslie suggests, then the prospects for meaningful change to be accomplished by adding quantifiers to our generalisations seem slim (cf. Saul, 2017: 14).

2 Semantics by opinion poll

I now turn to my argument that PTG is not a theory of the meaning of generics. While Leslie is not concerned with giving an account of the compositional semantics of generics, I take it that she does want to give a philosophically enlightening account of their meaning, since she begins one of her articles by puzzling over the truth-conditions of generics (Leslie, 2007: 375). But if this is what Leslie is aiming for, then asking only under what conditions ordinary speakers judge certain sentences to be true is not the right way to go about it.

Leslie’s favoured approach is founded on the understanding that “empirical work on what regulates people’s acceptance of generics can provide evidence in favor of one account of truth-conditions over another” (Lerner and Leslie, 2016: 406). While this generalisation is no doubt true, ordinary speakers’ views about the truth values of sentences or utterance are mediated by a wide variety of factors, beyond

⁶For example, Cheng (2011: 12) explicitly separates Leslie’s claims about what it takes for a generic to be true/false from her claims about the cognitive factors which underlie our dispositions to accept or reject generics. Cheng goes on to argue that the meaning of meaning sentences, such as ‘I mean *plus* by ‘+’’, should be associated with our dispositions to accept or reject them, and it is in the analysis of these dispositions, rather than the truth-conditions of meaning sentences, that PTG is invoked (ibid.: 15).

just their knowledge of truth-conditions. So, while this empirical data is certainly useful in testing hypotheses, it cannot be not the final word in an investigation into semantics.

The claims that Leslie makes in favour of PTG suggest that she takes speaker intuitions at face value. For example, she appeals to the judgments of a hypothetical speaker of English, whose response is “invariably immediate and confident” (Leslie, 2008: 1). Later, she says, “[i]f the relation between objective statistical frequencies and the truth-values of generics is in some way mediated by psychological considerations, it would be useful to determine exactly what role these considerations are playing” (ibid.: 16). Leslie goes on to suggest that the mediation of psychological considerations forms part of the truth-conditions of generics (or as she puts it, “how the world must be so as to conform to the claim” made by a generic (ibid.: 43)). But this is an incredibly strong claim; although it is in principle possible that such psychological considerations could enter the metaphysical picture, the more plausible hypothesis is that they mediate our truth value judgments without determining how the world must be in order for the sentences to be true.

2.1 Collecting truth value intuitions

When considering the empirical approach favoured by Leslie, we should ask how it is possible to pair speakers’ intuitions concerning the truth value of a particular generic and the conditions under which a speaker is making that judgment. In an ideal world, an experimental setting would fix the conditions and then elicit a truth value judgment. The truth value judgment would then be paired with the conditions. A representative example of this is the study in Schwarz (2013).⁷ In this study, which aimed to test a hypothesis concerning the semantics of plural definite descriptions, participants are asked to judge whether sentences with plural definite descriptions were true or false, based on an image. The image consisted of shapes of different colours displayed in an array and participants were asked to judge sentences such as ‘The circles were black’. By pairing judgments about sentences with particular states of affairs, e.g. when half the circles are black, it is possible to give a mean truth value judgment relative to a particular state of affairs.⁸

⁷For related examples of this experimental method in semantics, where the experimenter manipulates a carefully controlled context and asks their participants to make a judgment of some kind, see Križ and Chemla (2015) and Augurzky et al. (2023) on plural definite descriptions, Pietroski et al. (2009) on the interpretation of ‘most’, and Syrett, Bradley, et al. (2006), Syrett, Kennedy, et al. (2010), Aparicio et al. (2015), and Weicker and P. Schulz (2020) on gradable adjectives.

⁸Schwarz’s analysis is more nuanced than this; the experiment aimed to measure reaction times for truth value judgments relative to different states of affairs, in order to test a hypothesis

How would one go about performing such a task for generics? One problem with this approach is that generics are usually unbounded, in the sense that they resist contextual domain restriction (Krifka et al., 1995: 45; Drewery, 2005: 376-377; Sterken, 2015a: 13; although see Greenberg, 2007 for a revision of this position). This means that when we elicit truth value judgments for generics, we are asking speakers to make judgments based on their background conceptual knowledge, not on something controlled for by the experimenter.

Two individuals may therefore disagree over the truth value of a particular generic, not because they differ in their beliefs about its truth-conditions, but because they differ in their beliefs about whether those conditions obtain. For example, two people may disagree concerning the truth value of the following generic:

(6) Ships have sails.

One person may believe that a majority of ships have sails, while the other may believe that a majority of ships do not have sails. The first may therefore judge the sentence true, while the latter false, on the basis of this background information. Alternatively, one might hold that it is part of what it is to be a ship to have sails, while the other may believe that being a large boat is sufficient to be a ship, in which case again the first would judge the sentence true and the latter may judge it false. The two people might disagree on the truth value of (6) because they disagree about the nature of ships, even though, as in the first case, they agree on the truth-conditions of (6). So unlike in Schwarz's experiment described above, where the conditions against which sentences are judged true or false are strictly controlled, it is difficult to control for different background assumptions which may produce different truth value judgments even when participants agree on truth-conditions.

To illustrate this difficulty, consider an empirical study by Prasada, Khemlani, et al. (2013) which elicited truth value intuitions about a wide range of generics about familiar kinds from ordinary speakers. Participants were asked to judge the truth of sentences on a scale from Definitely True (+3) to Definitely False (-3). This study found evidence of very mixed truth value judgments for many generics. For example:

(7) Birds have wings. *Mean truth value judgment: 2.63*

(8) Lemons are sour. *Mean truth value judgment: 2.25*

That these apparently obviously true generics received mixed judgments suggests that participants evaluate them relative to different bodies of evidence. This is

concerning whether 'non-maximal' (i.e. exception-tolerating) readings of plural definite descriptions should be accounted for in their semantics or as a pragmatic phenomenon.

not surprising; conceptual representations are likely to be highly idiosyncratic and there is no reason to believe that one person’s representation of the concept BIRD or LEMON will be identical (or even could be identical) to another’s.

Another finding of this study was that participants were unwilling to call any generics completely false, even those often considered paradigm cases of false generics, such as ‘Books are paperbacks’ (0.21) or ‘Canadians are right-handed’ (−0.50) (Prasada, Khemlani, et al., 2013: Appendix B). Presumably, some of the generics that they tested are indeed completely false, for instance:

- (9) Plants are ferns. *Mean truth value judgment*: 0.04
- (10) Novels are mystery novels. *Mean truth value judgment*: 0.04

What are we to make of the result that these sentences were judged more true than false? First and foremost, I contend, it would be a mistake to take this data at face value and try to construct a semantic theory to vindicate it. We cannot tell from the data what conditions participants were evaluating these sentences relative to,⁹ so it is difficult if not impossible to draw conclusions about the truth-conditions of generics based on this data. It may be that participants attempted to accommodate background information in order to make the sentences more plausible, which would explain why so few generics were judged definitely false (the mean judgment for minority false generalisations like (9) and (10) was −0.52). This suggests that we must be very cautious when interpreting the results of this kind of study. It is not possible to fix the background conditions against which speakers make their truth value judgments about generics.

One method of countering this shortcoming is to use unfamiliar novel categories to elicit truth value judgments. For example, Cimpian, Brandone, et al. (2010) introduced participants to sentences about animals from a remote island, called “lorches”, and asked them to make judgments about sentences such as:

- (11) (a) Most lorches have purple feathers.
(b) Lorches have purple feathers.
- (12) (a) Most lorches have dangerous purple feathers.
(b) Lorches have dangerous purple feathers.

The results for generics were then compared with the sentences featuring the explicit quantifier ‘most’. Cimpian, Brandone, et al. showed that, amongst other things, unlike the (a)-sentences, participants were more likely to accept a generic when the property was dangerous than when it was not, for a given prevalence of the property among the category.

Admirable though this approach is, it is still not equivalent to the more controlled

⁹With two exceptions: participants were asked to make prevalence and cue validity estimates.

approach taken by Schwarz in the study concerning plural definite descriptions described above. We can assume that participants were using some information other than prevalence to make judgments about the truth of the sentences they were presented with. But beyond this, it is difficult to say precisely against what information the participants were using to assess the truth values of the generics, since we don't know what background assumptions the participants made about orches on the basis of the limited information they were given. Indeed there is a growing literature demonstrating that people form beliefs about the nature of novel categories described by generics on the basis of very little information (e.g. Gelman et al., 2010; Rhodes et al., 2012; Noyes and Keil, 2019; Vasilyeva and Lombrozo, 2020; Leshin et al., 2021; Benitez et al., 2022). So once again, even in an experimental setting, participants are making judgments about generics relative to opaque bodies of background knowledge, which cannot be controlled for. What we learn from this study is that whether a property is dangerous had an effect on truth value judgments, but we cannot say relative to what conditions participants were judging the sentences, so we cannot reach conclusions about people's beliefs about the truth-conditions of generics.

2.2 Top down semantics

This leads to another methodological question: when we present generics to ordinary speakers, why should we expect that they can reliably and accurately ascertain their truth values? Ordinary speakers are unlikely to attend to fine distinctions between semantic features of a sentence and non-semantic features which do not make a contribution to the truth-conditions of the sentence, but which affect judgments of truth.

To take unreflective truth value judgments as the starting point from which to develop a semantic theory, as Leslie does, is to take a 'top down' approach. This contrasts with a 'bottom up' approach on which we would proceed from the meanings of a sentence's constituents and composition rules to the truth-conditions of the sentence as a whole.¹⁰ The trouble with the top down approach is that it is likely to require a certain gerrymandering of the semantics to fit the data. This is because variations in intuitions may correspond to pragmatic or non-linguistic features of a sentence or utterance, as opposed to semantic features. These features can be subject to systematic error or inconsistency. As Kent Bach has remarked:

¹⁰Again, I should emphasise that Leslie explicitly denies that she is giving a compositional semantics for generics. As should be clear, I agree with this contention, but I differ from Leslie in my view of what her theory is really about.

Taking semantic intuitions seriously would make life miserable for semanticists, even more miserable than it already is. It would require doing semantics from the top down. . . . Covert constituents would have to be posited to provide the “residue” of meaning not accounted for by the overt ones. Or, alternatively, special meanings, hence ambiguity, would have to be attributed to certain of the overt constituents, insofar as their ordinary meanings seem not to make the right contribution to what is said. All this can be avoided if we don’t take people’s seemingly semantic intuitions too seriously (Bach, 2002: 26).

Leslie introduces both a covert constituent and ambiguity into her theory of the meaning of generics: first, she uses the covert operator, Gen, which plays an essential variable-binding function; second, she describes truth-conditions which vary radically based on the content of the generic.

These facts do not on their own discredit PTG as a semantic theory, but we should approach PTG with scepticism, because truth value intuitions should not form the basis of semantic theory when they are taken uncritically and at face value, rather than as data against which to test hypotheses, since they may be tracking non-semantic features of sentences. The best methodology for semantics represents a blend between the top down and bottom up approaches, where hypotheses about the contributions of individual words or morphemes and composition rules are assessed against data from truth value intuitions. Discrepancies can be resolved either by amending the hypotheses, or by proposing an error theory to explain away the discrepancy. But in PTG, the bottom up approach is completely absent, especially given the lack of an informative analysis of Gen. The purely top down approach leads to trouble in this particular case because the truth value judgments are so content-sensitive and unsystematic, which provides some at least *prima facie* evidence that there is something beyond their semantics mediating the judgments. Moreover, as I argued in the previous section, generics are distinctive in that is difficult to know against what conditions people are making truth value judgments, which makes those judgments especially bad guides to the truth-conditions of those sentences.

The complex and variable nature of Leslie’s truth-conditions for generics suggests that they are a result of “overfitting” the data of our intuitions (Williamson, 2020: 264): by failing to recognise that our evaluations of sentences do not perfectly track the truth values of those sentences, Leslie’s theory is made unnecessarily elaborate in its content-sensitivity. To put the complexity into the truth-conditions is to be overly credible in our treatment of speakers’ intuitions. Why think that ordinary speakers can simply intuit the truth values of sentences which concern the nature of our concepts and the relationships between them? Another, more plausible

explanation for the variability in truth value judgments is available: speakers use heuristics to evaluate the truth of statements that they are otherwise unable to verify.

2.3 Heuristics and truth value judgments

We can see the importance of not taking truth value intuitions too seriously elsewhere. A perplexing variability in truth value judgments would be observed for active/passive minimal pairs, since, as Kahneman notes, the grammatical subject makes different information accessible to the cognitive system (2003: 701):

- (13) Arsenal beat Tottenham.
- (14) Tottenham were beaten by Arsenal.

In these cases, judgments of acceptability do not track only semantic features of the pairs. But why would we assume that such divergence amounts to a difference in how the world must be in order for them to be true? To use them as evidence of truth-conditional divergence would be misguided (cf. Frege, 1960: §3).

Another helpful demonstration of this is the “inverse conjunction fallacy”, observed by Jönsson and Hampton (2006). This is a phenomenon, related to the conjunction fallacy observed by Tversky and Kahneman (1983), where people judge universal generalisations about subsets of a category to be less likely to be true than unmodified universal generalisations about the category as a whole. For instance, consider the pair:

- (15) All sofas have backrests.
- (16) All uncomfortable handmade sofas have backrests.

In Jönsson & Hampton’s studies, they found “that people have a strong tendency to accept unmodified universally quantified sentences [e.g. (15)] with greater frequency and/or greater confidence than the equivalent modified versions [e.g. (16)]” (2006: 328).¹¹ Given a standard account of the semantics of ‘all’ as a universal quantifier, (15) entails (16), since all uncomfortable handmade sofas are sofas. But we cannot recover this fact from experimental data of people’s truth value judgments, since some hold (15) to be true while (16) is false. If we started from that data, we would be led to posit a variable, highly content-sensitive semantics for ‘all’. We might even, in the same vein as Leslie, attempt to provide a complex account of what the world has to be like in order either (15) or (16) to be true which vindicates these intuitive judgments.

¹¹The same effect was found when the sentences were modified by ‘every single’ and ‘100% of’ (Experiment 4), which suggests that the effect is not linked to the word ‘all’.

This would be a mistake. An alternative and much more plausible route is to view the semantics of ‘all’ as simple and propose a hypothesis about human reasoning which can explain the discrepancy from what we would expect to see given the semantics.¹² For Jönsson & Hampton, this is a heuristic, which is fallible despite being generally useful in evaluating universal generalisations. When evaluating whether a universal generalisation is true or false, if we do not possess “an easily searchable exemplar space, we rely instead on the strength with which the predicate is represented as part of our concept” (2006: 331). In the case of the sofa, there is no easily searchable exemplar space of sofas (since we haven’t observed all the sofas in the world), so we default to representativeness to decide if the generalisation is true. This results in different judgments for the two sentences, since the representation or prototype of an uncomfortable handmade sofa is not as closely related with having a backrest as is the case for the prototypical sofa. We can in this way account for the inconsistency of the data without positing an inconsistent semantics for any of the components of the sentences. The result is essentially an error theory, which is both simpler than a PTG-style semantic theory and makes falsifiable predictions about the conditions under which people tend to make errors in their evaluation of the truth of universal generalisations.

Reflecting on this result, Hampton elaborates:

As philosophy and logic students are quickly made aware, universally quantified statements can only strictly be verified by an exhaustive search for counterexamples and the failure to find them. “*All swans are white*” is notoriously much harder to prove true than to prove false. Human language has opted therefore to take a different approach to determining “truth” of such statements, one which frequently leads to judgments that appear fallacious from a logical perspective. Statements of this kind are judged true on the basis of the strength and importance of their association with the content of a concept. To say that “*All S have property P*” in the context of everyday speech is to say something like “*An important and relevant fact about the kind S is that the property P is to be found in individuals of that kind*” (Hampton, 2012: 20).

Elsewhere, Hampton concludes: “a judgment of whether or not a property is true

¹²This follows from the following methodological principle, suggested by Sterken in her own discussion of PTG: “If a cognitive bias affects our intuitions in systematic ways, first try to treat these intuitions as mistaken, and only move to a semantic, metasemantic or pragmatic strategy if that fails” (2015b: 91). The general cognitive bias at work is presumably related to what Leslie calls the default mode of generalisation, but it seems likely on the basis of the evidence discussed here that its application is wider than just the evaluation of generics.

of a class may be made on the basis of how similar the class is to other classes for which the property is known to be generally true. We find it hard to differentiate our intuitions about what is generally true and what is always true” (2009: 97). Leslie herself seems willing to countenance such non-truth-conditional explanations in her own work on what she calls the Generic Overgeneralisation Effect, which is the tendency for people to recall generics as universal generalisations. Regarding Jönsson & Hampton’s experiments, Leslie concludes: “Either this constitutes a clear failure to appreciate the logical features of universal statements, or else their participants were interpreting these universals as generics” (Leslie, Khemlani, et al., 2011: 28). There is, quite rightly, no suggestion that the semantics of universal generalisations need to be revisited to account for the psychological evidence. So why not extend this view to the experimental data concerning generics? It is equally possible to take the variability in truth value judgments of generics as evidence that people do not ‘appreciate the logical features’ of generics, or that they default to another interpretation of the sentence in the absence of knowledge about all members of a kind.

I have shown that taking truth value judgments at face value would lead us to propose complex and variable truth-conditions for universal generalisations. Since this is a route that I suspect few working in semantics or cognitive science would like to take, I take it as good evidence that ordinary speakers’ truth value judgments can be mediated by heuristics, which can introduce inconsistency into our judgments.

2.4 Heuristics or meaning?

A defender of PTG’s status as a theory of the meaning of generics might object to the argument that our intuitions about the truth values of generics are mediated by heuristics by responding that they have an equally good explanation of why our intuitions about the truth values of generics are unsystematic. This is the Core Claim of PTG: the generalisations that generics express correspond to the judgments of the default mode of generalisation (Leslie, 2007: 381, 394), so the truth-conditions of generics correspond to the conditions under which our default mode of generalisation tells us to accept a generalisation (leaving space for empirical investigation into precisely what those conditions are). Since this mode of generalisation is unsystematic, our intuitions about the truth values of generics are too (ibid.: 397-398). This claim is independent of the specific clauses Leslie gives for the truth-conditions of different types of generics that I quoted above, so let’s focus just on the Core Claim.¹³

¹³I am grateful to an anonymous reviewer for raising this objection and suggesting this manner of presentation.

As a view about what the world has to be like in order for a generic to be true, the Core Claim of PTG is subject to two damaging objections. The first concerns what happens when two (or more) people's default modes of generalisation lead to different truth value intuitions for a particular generic (in a given context). Imagine that we are able to hold fixed the conditions against which those people evaluate the sentence. Which of these intuitions is 'correct'? If there is no way of deciding between these intuitions, then there seems to be no unique answer to the question of what the truth-conditions of that generic are. That means that if we pursue the cognitive explanation of the meaning of generics proposed by PTG, we have to choose between two unpalatable options: either each of our default modes of generalisation function in the same way (i.e. with the same inputs, they produce the same outputs), or the truth of generics is relativised to each distinct default mode of generalisation (i.e. each person). Even if the truth-conditions of generics are indeterminate or underdetermined in some way, they are presumably not different for each observer.

Second, even assuming unanimity in truth value judgments, with the clauses regarding the truth-conditions of generics removed, PTG would predict that any intuition about the truth value of a generic informed by the default mode of generalisation is 'correct'. Reduced to the Core Claim, then, PTG makes no falsifiable predictions, since if our default mode of generalisation leads us to endorse a generic, it is true. That result would be especially damaging for Leslie's project, as she wants to be able to maintain that striking property generics, like 'Mosquitoes carry the West Nile virus', are true, while superficially similar striking property generics about social groups, like 'Muslims are terrorists', are false (Leslie, 2017: 403). This illustrates the problem with this mode of overfitting the data: if PTG makes no falsifiable claims about the meaning of generics, it is not a good theory of the meaning of generics.

I suspect Leslie is aware of that flaw, which explains why she rightly appeals to evidence from the psychological literature so that she is able to make more systematic claims about our judgments of the truth of generics (Leslie, 2008: 16). Without these claims about the patterns in our judgments, PTG is vacuous. With the addition of the clauses stating the truth-conditions of different types of generics, the theory is at least falsifiable, but this does not protect PTG from the first criticism. By tying the truth-conditions of generics so closely to our judgments about the truth of generics, PTG leaves no room for disagreement and hence errors in our evaluations of generics. I therefore conclude that the Core Claim of PTG does not constitute a viable basis for a theory of the meaning of generics.

3 Evaluation Heuristics

In what follows, I will argue that the observation that fallible heuristics can mediate truth value judgments helps us to make sense of the variability in our judgments about the truth values of generics, and that PTG is fruitfully interpreted as a theory of these heuristics of judgment, not of the meaning of generics.

Heuristics of judgment are methods of making decisions under conditions of uncertainty or limited resources: “A heuristic is a strategy that ignores part of the information, with the goal of making decisions more quickly, frugally, and/or accurately than more complex methods” (Gigerenzer and Gaissmaier, 2011: 454). In practice, most decisions are made with limited cognitive resources or incomplete knowledge, so we frequently employ heuristics to help us make judgments.

3.1 Why heuristics?

First, one might ask why we would need to use heuristics to judge the truth of generics. The most convincing answer is that it is hard for us to know for certain whether a generic is true (Prasada, 2000: 67). We frequently need to make decisions in cases where there is either insufficient information or where there is no optimal solution to a problem and it is in these cases that the use of heuristics is appropriate (and perhaps even required). Judging the truth value of a generic may be just like this, since it is clear that even on reflection we (both ordinary speakers and semanticists) do not know what its truth-conditions are, let alone whether they obtain.

The truth-conditions could, for example, include some modal element (Asher and Morreau, 1995; Nickel, 2016; Thakral, 2018; Kirkpatrick, 2023), require the existence of some sort of pattern which must be recognised over time (Carlson, 2008; Cohen, 1999),¹⁴ rest on a notion of kind-predication where unobservable kinds inherit some of the properties of their members (Liebesman, 2011; Teichman, 2023), or they may even be context-sensitive (Sterken, 2015a; Lee and Nguyen, 2022), semantically incomplete (Nguyen, 2020), or unspecific with respect to the generalisations they express (Bosse, 2021). If these suggestions are along the right lines, it would be hard and maybe even impossible to verify a generic. But even if generics are just unbounded universals concerning all or almost all actual instances of a property, in practice we would not come to endorse or reject them by considering each instance in turn. Just as Hampton suggested in the passages quoted above,

¹⁴Van Rooij & Schulz (2020) explicitly draw a connection between Cohen’s probability theory of generics and Kahneman’s Heuristics and Biases research programme. The basis of this connection is the fact that people often mistake typicality for probability, i.e. if my representation of a typical F has the property G, I will consider it probable that an arbitrary F is G.

we do not in practice evaluate universal generalisations on the basis of the observation of every instance, but instead evaluate them on the basis of “the importance of their association with the content of a concept” (2012: 20).

To see some evidence for this claim, consider the ‘modifier effect’, which has been observed by several studies (e.g. Connolly et al., 2007; Hampton et al., 2011; Jönsson and Hampton, 2012). Participants were asked to rank the following sentences for likelihood to be true (Connolly et al., 2007):

- (17) Candles are made of wax (unmodified).
- (18) Scented candles are made of wax (prototypical modifier).
- (19) Purple candles are made of wax (non-prototypical modifier).
- (20) Expensive purple candles are made of wax (additional non-prototypical modifier).

The mean ranking across a set of target sentences decreased from (17) to (20), i.e. participants judged each sentence to be less likely to be true than the last. In an unpublished study reported in Hampton (2012), Hampton and Lam investigated the interactions between the modifier effect and the inverse conjunction fallacy. They found that the addition of ‘all’ into a bare plural sentence led, as one would expect, to a lower mean average in the judgments of the likelihood of the sentence being true. But the modifier effect was still significant for both the ‘all’-sentences and the bare plurals, although the effect was larger for bare plurals. Plausibly, then, the same heuristics are being used to evaluate both universal generalisations and generics, with the difference that for universal generalisations, people are less tolerant (although not completely intolerant) of counterinstances.

The fact that the universal generalisations received a lower mean rating than generics is consistent with the view that ‘all’ has a “maximizing” effect (Dowty, 1987). This is the view that ‘all’ has the effect of ruling out the possibility of exceptions to the generalisation (cf. Brisson, 2003; Križ, 2016). We do, of course, still exaggerate, so the effect of ‘all’ may in practice be only to rule out almost all, rather than all, exceptions (Jönsson and Hampton, 2006: 331). Nevertheless, the evidence supports the conclusion that we use heuristics to evaluate both generics and universal generalisations. These heuristics may lead us into inconsistency and do not necessarily reflect the truth-conditions of generics or universal generalisations.

Acknowledging the impossibility of its justification on the basis of deductive inference, Frank Ramsey described the practice of induction as a “habit of the mind” without which we would be “very much worse off” (1990: 92-93). We should think of the heuristics that guide our judgments of the truth of generics in much the same way; without these heuristics, we would be unable to judge whether or not

generics are true and, given the important role that they play in our conceptual knowledge (Prasada and Dillingham, 2006) and early education about the world (Cimpian and Markman, 2009), there is a clear utility in being able to make these judgments, even if they are fallible. Perhaps this is the best that we can do.

3.2 What is preserved?

Once we view PTG as a theory of how heuristics affect our judgements of the truth values of generics rather than of their truth-conditions, we are better equipped to understand some of their puzzling features.

3.2.1 Content-sensitivity

One of the challenges for giving a semantics for generics is that our truth value judgments appear to be sensitive to the wording and content of generics, which is what motivates Leslie's complex and variable truth specifications. Since heuristics can be sensitive to quite fine-grained distinctions in language (such as word order or choice), their use would explain why our truth value judgments appear very sensitive to the content of generics. Making things more concrete, we can modify Leslie's truth-conditions to suggest heuristics for generics.

Judge a generic, 'Ks are F' true when the following conditions hold:

- (21) (a) The counterinstances are negative.
- (b) If F is a property of a prototypical K then some Ks are F.
- (c) If K is an artifact or social kind then F is the function or purpose of the kind K.
- (d) If F is striking or dangerous then some Ks are F and the other Ks are disposed to be F.
- (e) If none of the antecedents of the three preceding conditions are true, then almost all Ks are F.

I do not claim that these heuristics correctly predict our truth value judgments for generics, merely that PTG can be restated as a theory of truth value judgments rather than worldly truth specifications.

The explanation for these heuristic rules will presumably relate to Leslie's default mode of generalisation; people make judgments in this way because of features of human cognition. The difference is now that this mode of generalisation plays no role in the *semantics* of generics. Condition (21d) calls for special consideration though, as we might ask why striking or dangerous properties are singled out in this way, since prototypical Ks are typically counterinstances to such generics. I suggest that the explanation for this heuristic would be purely pragmatic, in the

broadest sense of the word: it is advantageous to treat Ks as if they were F when F is a property that poses some kind of threat (e.g. carrying the West Nile Virus) (compare with Leslie, 2017: 396). So, on the basis that it is not generally possible to identify which mosquitoes are carriers of a disease without carrying out some kind of impractical testing, I might avoid all mosquitoes on the basis of the belief that some mosquitoes carry the West Nile Virus, and would utter or endorse the generic, ‘mosquitoes carry the West Nile Virus’. As Restall puts it, in these cases “it is much better to err on the side of *false positives* than *false negatives*” (2018: 2).

3.2.2 Inferential asymmetry

The inferential asymmetry of generics, the tendency to infer more from a generic than we require to believe that it is true (Cimpian, Brandone, et al., 2010), plays an important role in Leslie’s more recent work on the harmful effects of social generics (e.g. Leslie, 2017). Acknowledging the role of heuristics in our judgments of generics, we can now see that the inferential asymmetry of generics might be caused by the fact that we habitually accept false generics as true and then use these false generics to draw inferences, which are unsound because they have a false premise. If the conditions under which we accept generics as true bear no direct relation to prevalence, but instead to representativeness, as suggested by the model of bounded rationality presented in Tversky and Kahneman (1982), it is no surprise that there is a mismatch between the conditions under which we accept them as true and the extent of the inferences we use them to license.

This raises the question, why should the asymmetry hold in the direction that it does, i.e. why do we accept generics on the basis of little evidence but use them to license strong inferences, rather than the other way around? To give a satisfactory answer to this question would take me too far into the realms of cognitive psychology, but it may be the case people habitually overestimate the extent to which the properties of a prototypical instance of a category are prevalent among other instances of the category (cf. n. 14). This would be, as Cimpian, Brandone, et al. suggest (2010: 1475-1476), an exemplification of the tension between intensional reasoning (reasoning about a concept) and extensional reasoning (reasoning about objects which fall under a concept), since people may take a property to be typical of a concept even if it is not possessed by a high proportion of objects which fall under the concept. In the long run, it may be better to overgeneralise on the basis of a few instances than to undergeneralise. We are able, then, to make more and more confident predictions about new instances of categories that we encounter. But this does not, however, force us to posit any inconsistency in the semantics of generics. The inconsistency is in our heads.

3.2.3 Stereotyping and prejudice

As we have seen, the world contains too much information for us to process it all in detail, so our cognitive system must be selective in what information it chooses to process. Heuristics help us to navigate the complexity of the world, but at the same time they can lead us to make or endorse morally questionable generalisations when they concern social kinds, such as ethnic, religious or gender groups, as Leslie has argued (2014; 2017). This coheres with the view in social psychology that stereotyping is closely related to the use of heuristics of judgment (e.g. Rothbart et al., 1978; Bodenhausen, 1990);¹⁵ if we make judgments about generics using heuristics, it is no surprise that generics are associated with stereotyping since the same mechanism is responsible for our endorsement of both generics and stereotypes. This suggests that there is still mileage in the increasingly popular research programme concerning the links between the use of generics and prejudice.¹⁶

In their landmark discussion of the semantics of generics, Krifka et al. rejected the theory that generics express stereotypes, which they attribute to Declerck (1986) and Geurts (1985),¹⁷ on the basis that this ties the truth-conditions of generics too closely to subjective issues such as cultural norms (Krifka et al., 1995: 49); for example, the following sentence is false, regardless of whether it is held to be an item of cultural knowledge within a given community:

(22) Snakes are slimy.

Their conclusion is surely correct; generics are about the world and we can be mistaken about their truth values if we are mistaken about the nature of world. But is it just a coincidence that generics are closely associated with stereotyping? By acknowledging the role of heuristics in the evaluation of generics, we can see how the same psychological processes that are involved in stereotype formation are involved in the judgment of the truth values of generics.

Since on the model I have argued for here, universal generalisations are also judged using heuristics, one might ask why they are not equally associated with stereotyping. The answer is simple: universal generalisations are less tolerant of counterexamples than generics are, and so are more easily falsified (Lemeire, 2021). As such, it is much easier to show someone categorically that their belief in a universal

¹⁵See also Allport's description, in his seminal text, of prejudice as the result of the "normal and natural tendency to form generalizations, concepts, categories whose content represents an oversimplification of [our] world of experience" (1954: 27).

¹⁶On this issue, see Haslanger (2012), Langton et al. (2012), Leslie (2014), Leslie (2017), Wodak, Leslie, and Rhodes (2015), Wodak and Leslie (2017), Lemeire (2021), McKeever and Sterken (2021), and Bosse (2022).

¹⁷Recently, Bosse has reversed the explanatory direction and argued that generics express stereotypes precisely because the content of stereotypical beliefs is generic.

generalisation is false than it is for a generic, so it is easier to hold on to a belief in a generic in the face of counterevidence.

Reconceptualising PTG as a theory of heuristics can help us see how what is usually an effective and efficient way of thinking can lead us astray in certain cases, with potentially harmful effects (see also van Rooij and K. Schulz, 2020 for a similar suggestion). But, given that the problem is not located in the semantics of generics, this view encourages us to ask whether the focus on *generics* in particular is justified, since it points towards a more general psychological explanation, rather than one associated with a particular form of words.

4 Conclusion

I have argued that we should reconceptualise Leslie’s theory of the meaning of generics as a theory of how heuristics affect our judgements of the truth values of generics. Viewing the theory in this way sheds light on the link between generics and cognitive biases, and opens up the possibility of drawing connections between our disposition to judge certain generics as true and other cases of bias resulting from the employment heuristics of judgment. As PTG is compatible with a range of theories of the semantics of generics, it offers an opportunity to explain away some of the truth value intuitions which diverge from the predictions of our theories, making the daunting project of giving a semantics for generics just a little easier.

In uncovering and probing these heuristics, Leslie makes a valuable contribution to our understanding of generics. But it is important that philosophers, linguists and cognitive scientists keep in mind that PTG does not provide an account of the meaning of generics so that they do not overestimate the potential to change how humans think by changing how they speak.

References

- Allport, G. W. (1954). *The Nature of Prejudice*. Cambridge, Mass: Addison-Wesley.
- Aparicio, H., Xiang, M., and Kennedy, C. (2015). “Processing Gradable Adjectives in Context: A Visual World Study”. In: *Semantics and Linguistic Theory*, pp. 413–432. DOI: 10.3765/salt.v25i0.3128.
- Asher, N. and Morreau, M. (1995). “What Some Generic Sentences Mean”. In: *The Generic Book*. Ed. by G. N. Carlson and F. J. Pelletier. Chicago ; London: University of Chicago Press, pp. 300–338.
- Augurzky, P. et al. (2023). “Putting Plural Definites into Context”. In: *Proceedings of Sinn und Bedeutung* 27, pp. 19–32. DOI: 10.18148/sub/2023.v27.1050.

- Bach, K. (2002). “Seemingly Semantic Intuitions”. In: *Meaning and Truth: Investigations in Philosophical Semantics*. Ed. by J. K. Campbell, M. O’Rourke, and D. Shier. New York: Seven Bridges Press, pp. 21–33.
- Benitez, J., Leshin, R. A., and Rhodes, M. (2022). “The Influence of Linguistic Form and Causal Explanations on the Development of Social Essentialism”. In: *Cognition* 229, p. 105246. DOI: 10.1016/j.cognition.2022.105246.
- Bodenhausen, G. V. (1990). “Stereotypes as Judgmental Heuristics: Evidence of Circadian Variations in Discrimination”. In: *Psychological Science* 1.5, pp. 319–322. DOI: 10.1111/j.1467-9280.1990.tb00226.x.
- Bosse, A. (2021). “Generics: Some (Non) Specifics”. In: *Synthese* 199.5, pp. 14383–14401. DOI: 10.1007/s11229-021-03426-8.
- (2022). “Stereotyping and Generics”. In: *Inquiry*, pp. 1–17.
- Brisson, C. (2003). “Plurals, ‘all’, and the Nonuniformity of Collective Predication”. In: *Linguistics and Philosophy* 26.2, pp. 129–184.
- Carlson, G. N. (2008). “Patterns in the Semantics of Generic Sentences”. In: *Time and Modality*. Ed. by J. Guéron and J. Lecarme. Dordrecht: Springer Netherlands, pp. 17–38.
- Cheng, K.-Y. (2011). “A New Look at the Problem of Rule-Following: A Generic Perspective”. In: *Philosophical Studies* 155.1, pp. 1–21. DOI: 10.1007/s11098-010-9561-9.
- Cimpian, A., Brandone, A. C., and Gelman, S. A. (2010). “Generic Statements Require Little Evidence for Acceptance but Have Powerful Implications”. In: *Cognitive Science* 34.8, pp. 1452–1482. DOI: 10.1111/j.1551-6709.2010.01126.x.
- Cimpian, A. and Markman, E. M. (2009). “Information Learned from Generic Language Becomes Central to Children’s Biological Concepts: Evidence from Their Open-Ended Explanations”. In: *Cognition* 113.1, pp. 14–25. DOI: 10.1016/j.cognition.2009.07.004.
- Claveau, F. and Girard, J. (2019). “Generic Generalizations in Science”. In: *Erkenntnis* 84.4 (4), pp. 839–859. DOI: 10.1007/s10670-018-9983-x.
- Cohen, A. (1999). “Generics, Frequency Adverbs, and Probability”. In: *Linguistics and Philosophy* 22.3, pp. 221–253. DOI: 10.1023/A:1005497727784.
- Connolly, A. C. et al. (2007). “Why Stereotypes Don’t Even Make Good Defaults”. In: *Cognition* 103.1, pp. 1–22. DOI: 10.1016/j.cognition.2006.02.005.
- Declerck, R. (1986). “The Manifold Interpretations of Generic Sentences”. In: *Lingua* 68.2-3, pp. 149–188. DOI: 10.1016/0024-3841(86)90002-1.
- DeJesus, J. M. et al. (2019). “Generic Language in Scientific Communication”. In: *Proceedings of the National Academy of Sciences* 116.37, pp. 18370–18377. DOI: 10.1073/pnas.1817706116.

- Dowty, D. R. (1987). “Collective Predicates, Distributive Predicates and ‘All’”. In: *Proceedings of the Third Eastern States Conference on Linguistics* 3. Ed. by F. Marshall, A. Miller, and Z.-s. Zhang, pp. 97–117.
- Drewery, A. (2005). “The Logical Form of Universal Generalizations”. In: *Australasian Journal of Philosophy* 83.3, pp. 373–393. DOI: 10.1080/00048400500191966.
- Frege, G. (1960). “Begriffsschrift”. In: *Translations from the Philosophical Writings of Gottlob Frege*. Trans. by P. T. Geach and M. Black. 2nd ed. Oxford: Basil Blackwell, pp. 1–20.
- Gelman, S. A., Ware, E. A., and Kleinberg, F. (2010). “Effects of Generic Language on Category Content and Structure”. In: *Cognitive Psychology* 61.3, pp. 273–301. DOI: 10.1016/j.cogpsych.2010.06.001.
- Geurts, B. (1985). “Generics”. In: *Journal of Semantics* 4.3, pp. 247–255. DOI: 10.1093/jos/4.3.247.
- Gigerenzer, G. and Gaissmaier, W. (2011). “Heuristic Decision Making”. In: *Annual Review of Psychology* 62.1, pp. 451–482. DOI: 10.1146/annurev-psych-120709-145346.
- Greenberg, Y. (2007). “Exceptions to Generics: Where Vagueness, Context Dependence and Modality Interact”. In: *Journal of Semantics* 24.2, pp. 131–167. DOI: 10.1093/jos/ffm002.
- Hampton, J. A. (2009). “Stability in Concepts and Evaluating the Truth of Generic Statements”. In: *Kinds, Things, and Stuff: Mass Terms and Generics*. Ed. by F. J. Pelletier. New York: Oxford University Press, pp. 80–99. DOI: 10.1093/acprof:oso/9780195382891.003.0005.
- (2012). “Generics as Reflecting Conceptual Knowledge”. In: *Recherches linguistiques de Vincennes* 41 (41), pp. 9–24. DOI: 10.4000/r1v.2036.
- Hampton, J. A., Passanisi, A., and Jönsson, M. L. (2011). “The Modifier Effect and Property Mutability”. In: *Journal of Memory and Language* 64.3, pp. 233–248. DOI: 10.1016/j.jml.2010.12.001.
- Haslanger, S. (2012). “Ideology, Generics, and Common Ground”. In: *Resisting Reality*. New York: Oxford University Press, pp. 446–478. DOI: 10.1093/acprof:oso/9780199892631.003.0017.
- Jönsson, M. L. and Hampton, J. A. (2006). “The Inverse Conjunction Fallacy”. In: *Journal of Memory and Language* 55.3, pp. 317–334. DOI: 10.1016/j.jml.2006.06.005.
- (2012). “The Modifier Effect in Within-Category Induction: Default Inheritance in Complex Noun Phrases”. In: *Language and Cognitive Processes* 27.1, pp. 90–116. DOI: 10.1080/01690965.2010.544107.
- Kahneman, D. (2003). “A Perspective on Judgment and Choice: Mapping Bounded Rationality”. In: *American Psychologist* 58, pp. 697–720. DOI: 10.1037/0003-066X.58.9.697.

- Kirkpatrick, J. R. (2023). “Generic Conjunctivitis”. In: *Linguistics and Philosophy* 46, pp. 379–428. DOI: 10.1007/s10988-022-09371-0.
- Krifka, M. et al. (1995). “Genericity: An Introduction”. In: *The Generic Book*. Ed. by G. N. Carlson and F. J. Pelletier. Chicago ; London: University of Chicago Press, pp. 1–124.
- Križ, M. (2016). “Homogeneity, Non-Maximality, and ‘All’”. In: *Journal of Semantics* 33.3, pp. 493–539. DOI: 10.1093/jos/ffv006.
- Križ, M. and Chemla, E. (2015). “Two Methods to Find Truth-Value Gaps and Their Application to the Projection Problem of Homogeneity”. In: *Natural Language Semantics* 23.3, pp. 205–248. DOI: 10.1007/s11050-015-9114-z.
- Langton, R., Haslanger, S., and Anderson, L. (2012). “Language and Race”. In: *Routledge Companion to Philosophy of Language*. Ed. by G. Russell and D. G. Fara. London: Taylor & Francis Group.
- Lee, J. and Nguyen, A. (2022). “What’s Positive and Negative about Generics: A Constrained Indexical Approach”. In: *Philosophical Studies* 179.5, pp. 1739–1761. DOI: 10.1007/s11098-021-01727-9.
- Lemeire, O. (2021). “Falsifying Generic Stereotypes”. In: *Philosophical Studies* 178.7, pp. 2293–2312. DOI: 10.1007/s11098-020-01555-3.
- Lerner, A. and Leslie, S.-J. (2016). “Generics and Experimental Philosophy”. In: *A Companion to Experimental Philosophy*. Ed. by W. Buckwalter and J. Sytsma. Oxford: Wiley-Blackwell, pp. 404–416.
- Leshin, R. A., Leslie, S.-J., and Rhodes, M. (2021). “Does It Matter How We Speak About Social Kinds? A Large, Preregistered, Online Experimental Study of How Language Shapes the Development of Essentialist Beliefs”. In: *Child Development* 92.4, e531–e547. DOI: 10.1111/cdev.13527.
- Leslie, S.-J. (2007). “Generics and the Structure of the Mind”. In: *Philosophical Perspectives* 21, pp. 375–403.
- (2008). “Generics: Cognition and Acquisition”. In: *The Philosophical Review* 117.1, pp. 1–47. DOI: 10.1215/00318108-2007-023.
- (2014). “Carving Up the Social World with Generics”. In: *Oxford Studies in Experimental Philosophy*. Ed. by J. Knobe, T. Lombrozo, and S. Nichols. Oxford: Oxford University Press, pp. 208–230. DOI: 10.1093/acprof:oso/9780198718765.003.0009.
- (2017). “The Original Sin of Cognition: Fear, Prejudice, and Generalization”. In: *The Journal of Philosophy* 114.8, pp. 393–421. DOI: 10.5840/jphil12017114828.
- Leslie, S.-J., Khemlani, S., and Glucksberg, S. (2011). “Do All Ducks Lay Eggs? The Generic Overgeneralization Effect”. In: *Journal of Memory and Language* 65.1, pp. 15–31. DOI: 10.1016/j.jml.2010.12.005.

- Leslie, S.-J. and Lerner, A. (2022). *Generic Generalizations*. In: *The Stanford Encyclopedia of Philosophy*. Ed. by E. N. Zalta and U. Nodelman. Fall 2022. Metaphysics Research Lab, Stanford University.
- Liebman, D. (2011). “Simple Generics”. In: *Nous* 45.3, pp. 409–442. DOI: 10.1111/j.1468-0068.2010.00774.x.
- Mari, A., Beyssade, C., and Del Prete, F. (2013). “Introduction”. In: *Genericity*. Ed. by A. Mari, C. Beyssade, and F. Del Prete. 43. Oxford: Oxford University Press. DOI: 10.1093/acprof:oso/9780199691807.003.0001.
- McKeever, M. and Sterken, R. K. (2021). “Social and Political Aspects of Generic Language and Speech”. In: *The Routledge Handbook of Social and Political Philosophy of Language*. Ed. by R. K. Sterken and J. Khoo. Routledge, pp. 259–280. DOI: 10.4324/9781003164869-20.
- Nguyen, A. (2020). “The Radical Account of Bare Plural Generics”. In: *Philosophical Studies* 177.5, pp. 1303–1331. DOI: 10.1007/s11098-019-01254-8.
- Nickel, B. (2016). *Between Logic and the World: An Integrated Theory of Generics*. Oxford: Oxford University Press. DOI: 10.1093/acprof:oso/9780199640003.001.0001.
- (2017). “Generics”. In: *A Companion to the Philosophy of Language*. Ed. by B. Hale, C. Wright, and A. Miller. John Wiley & Sons, Ltd, pp. 437–462. DOI: 10.1002/9781118972090.ch18.
- Noyes, A. and Keil, F. C. (2019). “Generics Designate Kinds but Not Always Essences”. In: *Proceedings of the National Academy of Sciences of the United States of America* 116.41, pp. 20354–20359.
- Peters, U. (2021). “Science Communication and the Problematic Impact of Descriptive Norms”. In: *The British Journal for the Philosophy of Science*. DOI: 10.1086/715001.
- Pietroski, P. et al. (2009). “The Meaning of ‘Most’: Semantics, Numerosity and Psychology”. In: *Mind & Language* 24.5, pp. 554–585. DOI: 10.1111/j.1468-0017.2009.01374.x.
- Prasada, S. (2000). “Acquiring Generic Knowledge”. In: *Trends in Cognitive Sciences* 4.2, pp. 66–72. DOI: 10.1016/S1364-6613(99)01429-1.
- Prasada, S. and Dillingham, E. M. (2006). “Principled and Statistical Connections in Common Sense Conception”. In: *Cognition* 99.1, pp. 73–112. DOI: 10.1016/j.cognition.2005.01.003.
- Prasada, S., Khemlani, S., et al. (2013). “Conceptual Distinctions amongst Generics”. In: *Cognition* 126.3, pp. 405–422. DOI: 10.1016/j.cognition.2012.11.010.
- Ramsey, F. P. (1990). “Truth and Probability”. In: *Philosophical Papers*. Ed. by D. H. Mellor. Cambridge: Cambridge University Press, pp. 52–109.
- Restall, G. (2018). “Accommodation, Inference, Generics and Pejoratives”. In.

- Rhodes, M., Leslie, S.-J., and Tworek, C. M. (2012). “Cultural Transmission of Social Essentialism”. In: *Proceedings of the National Academy of Sciences* 109.34, pp. 13526–13531. DOI: 10.1073/pnas.1208951109.
- Rosola, M. and Cella, F. (2020). “Generics and Epistemic Injustice”. In: *Ethical Theory and Moral Practice* 23.5, pp. 739–754. DOI: 10.1007/s10677-020-10095-y.
- Rothbart, M. et al. (1978). “From Individual to Group Impressions: Availability Heuristics in Stereotype Formation”. In: *Journal of Experimental Social Psychology* 14.3, pp. 237–255. DOI: 10.1016/0022-1031(78)90013-6.
- Saul, J. (2017). “Are Generics Especially Pernicious?” In: *Inquiry*, pp. 1682–1699. DOI: 10.1080/0020174X.2017.1285995.
- Schiller, H. I. (2023). “Genericity and Inductive Inference”. In: *Philosophy of Science* 90.3, pp. 539–557. DOI: 10.1017/psa.2022.94.
- Schwarz, F. (2013). “Maximality and Definite Plurals - Experimental Evidence”. In: *Proceedings of Sinn und Bedeutung* 17, pp. 509–526.
- Sterken, R. K. (2015a). “Generics in Context”. In: *Philosopher’s Imprint* 15.21, pp. 1–30.
- (2015b). “Generics, Content and Cognitive Bias”. In: *Analytic Philosophy* 56.1, pp. 75–93. DOI: 10.1111/phib.12056.
- (2015c). “Leslie on Generics”. In: *Philosophical Studies* 172.9, pp. 2493–2512.
- (2017). “The Meaning of Generics”. In: *Philosophy Compass* 12.8, e12431. DOI: 10.1111/phc3.12431.
- Syrett, K., Bradley, E., et al. (2006). “Shifting Standards: Children’s Understanding of Gradable Adjectives”. In: *Proceedings of the Inaugural Conference on Generative Approaches to Language Acquisition — North America*. Vol. 2. Cambridge, MA: UConn Occasional Papers in Linguistics 4, pp. 353–364. DOI: 10.7282/T39W0HHS.
- Syrett, K., Kennedy, C., and Lidz, J. (2010). “Meaning and Context in Children’s Understanding of Gradable Adjectives”. In: *Journal of Semantics* 27.1, pp. 1–35. DOI: 10.1093/jos/ffp011.
- Teichman, M. (2023). “The Sophisticated Kind Theory”. In: *Inquiry* 66.9, pp. 1613–1654. DOI: 10.1080/0020174X.2016.1267407.
- Thakral, R. (2018). “Generics and Weak Necessity”. In: *Inquiry*, pp. 1–28. DOI: 10.1080/0020174X.2018.1426683.
- Tversky, A. and Kahneman, D. (1982). “Availability: A Heuristic for Judging Frequency and Probability”. In: *Judgment under Uncertainty: Heuristics and Biases*. Ed. by A. Tversky, D. Kahneman, and P. Slovic. Cambridge: Cambridge University Press, pp. 163–178. DOI: 10.1017/CB09780511809477.012.

- (1983). “Extensional versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment”. In: *Psychological Review* 90.4, pp. 293–315. DOI: 10.1037/0033-295X.90.4.293.
- Van Rooij, R. and Schulz, K. (2020). “Generics and Typicality: A Bounded Rationality Approach”. In: *Linguistics and Philosophy* 43.1, pp. 83–117. DOI: 10.1007/s10988-019-09265-8.
- Vasilyeva, N. and Lombrozo, T. (2020). “Structural Thinking about Social Categories: Evidence from Formal Explanations, Generics, and Generalization”. In: *Cognition* 204, p. 104383. DOI: 10.1016/j.cognition.2020.104383.
- Weicker, M. and Schulz, P. (2020). “Not All Gradable Adjectives Are Vague – Experimental Evidence from Adults and Children”. In: *Proceedings of Sinn und Bedeutung* 24.2 (2), pp. 406–422. DOI: 10.18148/sub/2020.v24i2.905.
- Williamson, T. (2020). *Suppose and Tell: The Semantics and Heuristics of Conditionals*. Oxford: Oxford University Press. DOI: 10.1093/oso/9780198860662.001.0001.
- Wodak, D. and Leslie, S.-J. (2017). “The Mark of the Plural: Generic Generalizations and Race”. In: *The Routledge Companion to Philosophy of Race*. Ed. by P. C. Taylor, L. M. Alcoff, and L. Anderson. New York: Routledge, pp. 277–289.
- Wodak, D., Leslie, S.-J., and Rhodes, M. (2015). “What a Loaded Generalization: Generics and Social Cognition”. In: *Philosophy Compass* 10.9, pp. 625–635. DOI: 10.1111/phc3.12250.