Keeping Context in Mind: A Non-Semantic Explanation of Apparent Context-Sensitivity

Abstract

Arguments for context-sensitivity are often based on judgments about the truth values of sentences: a sentence seems true in one context and false in another, so it is argued that the truth conditions of the sentence shift between these contexts. Such arguments rely on the assumption that our judgments reflect the actual truth values of sentences in context. Here, I present a non-semantic explanation of these judgments. In short, our judgments about the truth values of sentences are driven by heuristics that are only fallible reflections of actual truth values. These heuristics can lead to different truth-value judgments in different contexts, even when the sentence at issue is not semantically context-sensitive. As a case study, I consider Sterken's (2015a) argument for the context-sensitivity of generic generalisations. I provide a non-semantic explanation of Sterken's truth-value judgments, which builds on Leslie's (2007; 2008) theory of default generalisation.

Keywords: Generic generalisations, primitive generalisation, default generalisation, psychological heuristics, application conditions.

Compliance with Ethical Standards: No conflicting interests to disclose.

Acknowledgements: My thanks to Jake Quilty-Dunn, Andrei Cimpian, and Lawrence Barsalou for helpful literature recommendations. Thanks also

to Jiwon Kim, members of NYU's Cognitive Development Lab, Mirela Fus, Ravi Thakral, Oliver Lemeire, anonymous reviewers, and the editor for helpful comments and suggestions. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101063848.

1 Introduction

Arguments for context-sensitivity are often based on judgments about the truth values of sentences in context. The author of the argument invites us to agree with their judgment that some sentence has different truth values in different contexts, despite there being no significant change in the state of the world between those contexts. The author concludes that the truth conditions of the sentence vary with context.

This line of argument assumes that our judgments accurately reflect actual truth values. There is reason to think, however, that our judgments are often based on heuristics that only fallibly reflect actual truth values. This opens up an alternative explanation of the data that allegedly establishes context-sensitivity: our heuristics might lead to different truth-value judgments in different contexts, even when the sentence at issue is not semantically context-sensitive. This possibility will be illustrated by considering Rachel Sterken's (2015a) argument for the context-sensitivity of generic generalisations. Sterken's argument is a particularly good case study for two reasons. First, Sterken's argument is more nuanced than other arugments for context-sensitivity, relying not only on judgments about the truth values of generics in context but also on judgments about the truth values of other expressions. Second, heuristics for the assessment of truth values are generally an understudied phenomenon but Sarah-Jane Leslie (2007; 2008) has proposed a very clear series of heuristics for

the assessment of generic generalisations.

This paper is not intended to refute Sterken's argument for a context-sensitive account of generics but rather to show that it is incomplete, as there is an alternative explanation that she has not considered. A sentence may appear to change its truth value between contexts, not because its truth conditions change between contexts, but because our heuristics for assessing truth and falsity lead us to different answers in different contexts. More broadly, this paper presents a challenge for accounts of context-sensitivity in general. Authors of context-sensitivity arguments should either provide reasons to think that the relevant truth-value judgments do not stem from fallible heuristics, or they should give reasons to think that those heuristics are accurate in the cases at issue.

The following section will describe the distinction between truth conditions and heuristics in more detail, drawing on work by Johnston and Leslie (2012, 2019). It will be argued that truth-value judgments are fallible evidence for context-sensitivity because context might affect our truth-value judgments without affecting truth conditions. Section 3 will introduce Sterken's argument for the context-sensitivity of generics. Section 4 will then describe the heuristics that Leslie takes to determine our truth-value judgments about generic generalisations and 5 will explain how these truth-value judgments might be influenced by context.

2 Heuristics

2.1 Heuristics vs Truth Conditions

Some questions are difficult, or even impossible, to answer with complete reliability. To answer a question like that, we need to use a *heuristic*, i.e. "a simple procedure that helps find adequate, though often imperfect, answers to difficult

questions" (Kahneman, 2012, p. 98). How often do fires break out when their residents go on holiday? I don't know the answer to this question and don't have nearly enough information to answer this question through, say, a completely reliable algorithm. Instead, I have to rely on some heuristic to estimate the answer. Tversky and Kahneman (1974) identified several heuristics that are used to answer questions about probability and frequency. Using the *availability heuristic*, for example, your judgment will be based on the ease with which you can recall instances of similar fires; the easier it is, the more likely you will judge it to be. This heuristic is not completely unrelated to the question. If fires are rare, you'll likely have heard of very few and find them difficult to recall; if fires are common, you will probably find it easier. But the heuristic is fallible. Perhaps fires are generally rare but people close to you have been particularly unlucky.

Heuristics influence our thinking in vast numbers of judgments. Some heuristics are applied consciously, as when i estimate $589 + 10,059$ by adding 600 to $10,000$. Others are used unconsciously. When estimating probabilities, for example, you might not notice that your judgment is based on ease of recall. In this paper, I'll be focusing on the heuristics that we use to decide whether a sentence is true.

In discussing the heuristics that guide our truth-value judgements, I will extend a distinction made by Johnston and Leslie (2012, 2019) between heuristics and *application conditions*. Application conditions are the conditions under which a term applies. A theory of application conditions might tell us, for example, that the word 'gold' applies only to the element with atomic number 79. Such theories are often phrased in terms of the 'reference', 'semantic content', or 'extension' of the term.

Speakers often want to avoid using words whose application conditions are

not satisfied (excepting, e.g., cases of lies, jokes, metaphor etc.) but often have no way to tell whether these conditions are satisfied without relying on heuristics. Suppose I am describing a novel piece of jewelry and have to decide whether it counts as 'gold.' I don't have direct access to its chemical composition, or the expertise the assess it with certainty, so I have to rely on some fallible procedure for deciding whether the application conditions of 'gold' are satisfied. These heuristics might include colour, texture, shape, and hallmarks.[1] While these heuristics are good enough in normal circumstances, they are clearly fallible; a substance might satisfy all these superficial criteria and yet fail to be gold.

Theories of application conditions and theories of heuristics also differ in their relationship to linguistic behaviour.[2] Here's a very modest theory of the application conditions of 'bear': the term does not apply to dogs. Suppose someone says, in earnest, 'Watch out! There's a bear!' upon seeing a particularly large dog. They have used the term 'bear' with reference to a dog but this fact doesn't refute the modest theory above. It isn't the theory of application conditions that is wrong in this case, it is the speaker.[3] Of course, we shouldn't *always* conclude that a speaker is in error when their usage fails to match a theory of application conditions. Theories can be wrong too and language use might be a good source of evidence against a theory. The point is only that speaker error is a possible explanation of a mismatch between a theory of application conditions and linguistic behaviour. A theory of heuristics is more directly falsifiable by appeals to linguistic behaviour, however. If my theory

[1] In the psychology literature, which is usually concerned with heuristics rather than application conditions, the combination of these heuristic features can be referred to as my *concept* or *prototype* of gold. See (Knobe, 2003, pp. 314-5) for two different notions of a concept. For the classic pioneering work on prototypes, see (Rosch, 1973, 1978; Rosch and Mervis, 1975).

[2] See (Knobe, 2003, pp. 314-5) for this way of distinguishing application conditions from heuristics, though Knobe doesn't use these terms.

[3] There is debate about what kind of error this is. See Wikforss (2001) and Hattiangadi (2006) for arguments that application conditions are not prescriptive and Fennell (2013) for a response.

suggests that people will apply 'gold' under certain conditions, but they in fact do not, there is no possibility of impugning their usage; I am simply wrong about the operative heuristics.

I extend Johnston and Leslie's distinction from words to sentences. The application conditions of a sentence are its truth conditions. Just as the application conditions of a word like 'gold' tell us the conditions under which the word applies, the truth conditions of a sentence tell us the conditions under which the sentence applies. Just as the heuristics associated with a term are our ways of deciding whether its application conditions are satisfied, the heuristics associated with a sentence are our ways of deciding whether its truth conditions are satisfied. Suppose, for example, that my heuristic for identifying gold is sensitive only to its hallmarks. Then I will judge 'This is gold' to be true only when the referent of 'this' is stamped with the correct hallmarks. Clearly, however, being stamped with the correct hallmarks does not turn a copper-zinc alloy into gold. Though my heuristics for assessing the truth of the sentence are satisfied, the truth conditions of the sentence are not. The heuristics we use to assess the truth of the sentence are therefore distinct from its truth conditions.

2.2 Heuristics in Action

The distinction between application conditions and heuristics has historical antecedents. Putnam (1975), for example, argued that meanings (application conditions) are not in the head. There could be two individuals with identical psychologies who refer to different things by 'gold' due to their environments. In one's head are not meanings but, for Putnam, "stereotypes", which are "features which in normal situations constitute ways of recognizing if a thing belongs to the kind" (p. 147). Application conditions are the conditions under which a thing belongs to a kind and Putnam's stereotypes are the heuristics that we use

to decide whether application conditions are satisfied. Likewise, Kripke (1980) argued (although not in these terms) that the application conditions of proper names are not determined by the heuristics that we use to identify their bearers (Johnston and Leslie, 2019, p. 197). I might identify Gödel as the first person to prove the incompleteness of arithmetic, yet for all I know it is possible that Gödel in fact stole the proof, which shows that this heuristic is a fallible means of assessing whether the application conditions of 'Gödel' are satisfied.

The distinction between heuristics and application conditions opens up the study of language in a similar way to Grice's (1989) notion of implicature. Grice taught us that semantic theories are not totally at the mercy of truth-value intuitions because those intuitions might result from implicatures. Likewise, even when no implicature is present, truth-value intuitions might result from heuristics, rather than knowledge of application conditions. Some sentences might seem true or seem false, not because they are, but because of the defeasbile heuristics that we use to assess truth and falsity. Williamson (2020), for example, suggests that the truth conditions of the indicative conditional are captured by the material conditional but that various features of the indicative are explained by the fallible heuristics that we use to decide whether a conditional is true. Just as Grice was able to explain various intuitions about truth and falsity through general features of rationality, we can explain some of these intuitions through general features of cognition, namely the use of heuristics.

2.3 Heuristics and Context-Sensitivity

Arguments for context-sensitivity often rely on the judgment that the truth value of some sentence differs between contexts. As our truth-value judgments are the result of heuristics, however, they might be explained by the effect on context on those heuristics, rather than semantic context-sensitivity.

Context can clearly affect the outcome of a heuristic-driven process. We have already noted the availability heuristic, through which probability and frequency are estimated by ease of recall. Features of context that do not affect probability or prevalence can influence ease of recall. Gabrielcik and Fazio (1984), for example, presented participants with a questionnaire asking them to compare the frequency of two letters, e.g. "Do more words contain T or S?" Participants responded on a nine-point scale anchored to "Many more contain S" and "Many more contain T". They found that participants primed with words containing the letter T gave higher estimates for the relative frequency of that letter.

Gabrielcik and Fazio suggest that priming increases the availability of words beginning with T, which leads to higher estimates due to the availability heuristic. For our purposes, it doesn't matter which heuristic is being used. Whatever the heuristic at play, its result is affected by features of context. There is an implied linguistic result here: subjects primed with the letter T are more likely to judge that 'Slightly more words contain T' is true, but the priming is semantically irrelevant to the truth of the sentence; exposure to certain words does not affect the frequency with which letters appear in English. Priming effects have been demonstrated for a variety of tasks and a great many of those will extend into linguistic results, showing how semantically irrelevant features of context can influence truth-value judgments.

2.4 Summary

This section has introduced the distinction between application conditions and heuristics. The application conditions of a linguistic term tell us what the term applies to and the heuristics associated with a term are our fallible ways of deciding whether those application conditions are satisfied. The application con-

ditions of a sentence are its truth conditions and the heuristics associated with the sentence are our fallible ways of deciding whether the sentence is true. These heuristics can lead to different judgments in different contexts, which allows us to explain contextual variation in truth-value judgments without positing semantic context-sensitivity. The following section describes Sterken's indexical analysis of generics and the data she presents in favour of it. Sections 4 and 5 will then offer a non-semantic explanation of this data in terms of the heuristics by which we judge the truth values of generics.

3 Generic Generalisations

3.1 Generics and Context-Sensitivity

If I tell you that some contemporary philosophers are boring, you know how to tell whether I speak the truth: go look for a boring philosopher. If you find one, then I have spoken the truth. If I tell you that all contemporary philosophers are boring, you can follow the same procedure: go check the philosophers. If each of them is boring then I have spoken truly. Other generalisations are more complex. If I tell you that 'many' contemporary philosophers are boring, for example, then you have to decide what counts as 'many'. Once this issue is settled, however, you can follow the same procedure: go and check the philosophers.

These generalisations all tell us something about the world, so we can go and check the world to figure out whether they have told us something true. What if I say simply that contemporary philosophers are boring? How would you verify that claim? The problem is vexed because this generalisation includes no quantifier (like 'some', 'all', or 'many') that tells you how far the generalisation extends. This is the mark of what are called *generic generalisations* or simply *generics*.[4] It is usually assumed that a generic, like a quantified generalisation,

[4]Here, I restrict discussion to bare plural generics. Interesting complications are raised by,

tells us something that is true under certain conditions. Research has then focused on trying to identify these truth conditions. But, despite the prevalence of generics in natural language, there remains no satisfactory account of the conditions under which they are true or false.[5]

Sterken (2015a) points to a phenomenon that makes the truth conditional analysis of generics still more difficult: a single generic can seem to have different truth values in different contexts. Consider, for example:

Indians: Indians eat beef.[6]

The truth value of this generic seems to vary between contexts, even when Indian beef consumption is held constant. Consider, for example:

Context 1: Beef is popular all over the world. Europeans eat beef. They eat beef in South America. Indians eat beef.

Context 2: Despite the cultural taboo prevalent in India, Indians eat beef, but it is less common than in Europe.

Indians seems false in Context 1 but same generic seems true in Context 2.

3.2 Sterken's Indexial Account

The standard syntactic analysis takes generics to be structurally identical to adverbially-quantified generalisations like 'Contemporary philosophers are mostly boring,' which are analysed as exhibiting the form:

Mostly [Contemporary philosopher x] [Boring x]

say, definite and indefinite generics. See Greenberg (2007) for discussion of indefinite singulars.

[5] See Leslie (2008; 2007) and Sterken (2015b) for problems with some of the most influential accounts to date.

[6] The example is inspired by one in Cohen (2004).

The 'Mostly' operator binds the variables, resulting in a sentence that is true just in case most variable assignments that satisfy 'Contemporary philosopher' also satisfy 'Boring'.[7] The generic 'Contemporary philosophers are boring' is then analysed as:

Gen [Contemporary philosopher x] [Boring x].

Here, the variable-binding role is played by an unpronounced operator *Gen*. The question then is how *Gen* determines the truth conditions of the generic as a function of the explicitly pronounced words.[8]

Sterken argues that the context-sensitivity of generics is best explained by analysing Gen as an indexical quantifier. The paradigm indexicals are words like 'I', 'here', and 'now', which vary in reference depending on the context in which they are used. 'I am hungry', for example, is true when uttered by a hungry person but false when uttered by a non-hungry person because 'I' refers to different people in different contexts.

Sterken suggests that the indexical quantifier Gen is context-sensitive in two ways. The context must set the quantificational force of the operator as well as what Sterken calls the 'lexical restrictor'. 'Sometimes' for example has existential force (like 'some') and is restricted to *actual* situations. In contrast, 'Normally' plausibly has universal force (like 'all') but is restricted to *normal* situations that may not be actual. Sterken's account allows for the lexical restrictor and quantificational force of Gen to vary independently, creating complex interpretations that might not be easy to pin down using natural language quantifiers.

As an example of variation in the lexical restrictor, consider:

[7]This analysis of adverbially-quantified generalisations was developed by Heim (1982), expanding on Lewis (1975). Kamp (1981/84) independently developed a theory very similar to Heim's.

[8]The idea that generics involve an unpronounced binding operator is suggested by Heim (1982, pp. 127-128).

Post: Workers in Sorting Room 6 handle the mail from Antarctica.[9]

Suppose that there has never been any mail from Antarctica but that the mail system is set up in such a way that Sorting Room 6 would deal with mail from Antarctica, were it to arrive. Now consider the following two linguistic contexts:

Context 1: What do they do in Sorting Room 6 do all day? Workers in Sorting Room 6 handle the mail from Antarctica.

Context 2: The mail system is prepared for every eventuality. We have assigned people to handle mail from the unlikeliest of place. Workers in Sorting Room 6 handle the mail from Antarctica and workers in Sorting Room 7 handle the mail from Mars!

In Context 1, the generic seems false. No mail has ever arrived from Antarctica, so that certainly isn't keeping Sorting Room 6 busy. For the generic to be true in that context, there must be actual situations in which Sorting Room 6 handles mail from Antarctica. In Context 2, however, the generic seems true. Although no mail has ever arrived from Antarctica, Sorting Room 6 would handle that mail, were it ever to arrive. In this case, the generalisation is not restricted to actual situations but quantifies over possible situations as well.

As an example of variation in quantificational force, recall **Indians**. In Context 1, the level of beef-eating in India does not seem sufficient to make the generalisation true. In Context 2, however, the same level of beef consumption does seem sufficient, suggesting that the quantificational force has varied between these contexts.

[9]This is a variation of an example from Krifka et al. (1995, p. 72).

3.3 Sterken's Argument for the Indexical Account

Sterken's primary example of generic context-sensitivity, drawn from Nickel (2008), is:

Dobermans: Dobermans have floppy ears.

The ears of the Doberman are naturally floppy but are cut by breeders to give the pointy shape they are often associated with. Sterken notes that **Dobermans** is intuitively true when uttered in the context of evolutionary biology:

Context 1: Some breeds of dogs have evolved to focus on their hearing. These breeds have pointy ears. Dobermans, however, mostly rely on their sense of smell, which is why Dobermans have floppy ears.

This same generic is intuitively false when uttered in the context of a dog show in which all the Dobermans present have pointy ears, rendering the negation true:

Context 2: Welcome to this year's meeting of the Westminster Kennel Club. Once again, we've got a great range of dog appearances. While Labradors and golden retrievers have floppy ears, Dobermans do not. Dobermans have pointy ears.[10]

Sterken argues for the context-sensitivity of generics through a process of elimination. So what are the alternatives to be eliminated? How else might we explain the difference between Contexts 1 and 2? We might think that one of the pronounced words is responsible. Perhaps, in Context 1, 'Dobermans' is interpreted to mean something like 'Dobermans that haven't been tampered with by humans', while in Context 2, 'Dobermans' is interpreted to mean something like 'Dobermans at this dog show'. Or perhaps, in Context 1, 'have floppy

[10] This phrasing is taken from Nickel (2016).

ears' is interpreted to mean 'have floppy ears at birth', while in Context 2 it is interpreted to mean 'have floppy ears at this dog show'. Alternatively, we might think that our interpretation of the pronounced words is stable but that differences in truth value can be explained through standard features of quantifiers such as domain restriction. Gen might only quantify over Dobermans at the dog show in Context 2, for example, but be unrestricted in Context 1. Finally, we might think that the sentences itself is interpreted identically in both contexts but that some pragmatic phenomenon, like implicature or presupposition, leads to an apparent shift in truth value.

Sterken intends to discredit these suggestions through her A-Quantifier Test.

A-Quantifier Test: Check whether explicitly adverbially-quantified sentences vary their truth value across the same contexts as generics. If there is no difference in truth value, this is evidence that generics shift their truth values as a result of distinctive context-sensitivity in the semantics of the implicit quantifier Gen.[11]

Suppose that our interpretation of the subject or predicate were shifting between contexts. In that case, Sterken suggests, we would expect the same shift to occur for adverbially-quantified variations of the generic (Sterken, 2015b, p. 2505), e.g.:

A-Dobermans: Typically/Generally/Normally Dobermans have floppy ears.

It is true, for example, that typically Dobermans that have not been tampered with by humans have floppy ears but it is false (let us suppose) that typically Dobermans prepared for a dog show have floppy ears. According to Sterken, however, the quantified generalisations in **A-Dobermans** sound false in both

[11] The wording of the test is a synthesis of Sterken's three discussions of the test; two in Sterken (2015a) and one in Sterken (2015b).

contexts, showing that the interpretation of the subject term is not shifting in this way.

Sterken uses the same form of argument to suggest that standard features of quantification like domain restriction are not responsible for the shift in apparent truth value. If they were, we would expect to see it mirrored in adverbially-quantified generalisations (Sterken, 2015a, pp. 14-15). Likewise for pragmatic mechanisms like implicature or question-sensitivity.[12] Sterken assumes that generics of the form 'Ks are F' are "close in meaning" to adverbially-quantified generalisations like 'Typically/Generally/Normally Ks are F' (Sterken, 2015a, p. 10). If the generics shifted their truth values because of non-semantic factors like implicature, Sterken argues that those same factors would lead adverbially-quantified generalisations to shift their truth values relative to the same contexts.

Through this process of elimination, Sterken concludes that the context-sensitivity of generics is due to the semantics of the implicit Gen operator and that this operator displays a distinctive form of context-sensitivity that isn't shared by explicit quantifiers, opting ultimately for an analysis on which the operator is an indexical quantifier.

I find Sterken's judgment about **A-Dobermans** hard to accept. To my ear, the adverbially-quantified generalisation is false in the context of the dog show but true in the context of evolutionary biology. In the latter context, it seems natural to focus on what is *evolutionarily* normal, typical, or generally true, which excludes Dobermans that have suffered accidents, or have been tampered with by humans. Looking at the Dobermans that remain, they typically/generally/normally have floppy ears. I therefore find it hard to agree that the A-Quantifier Test, as applied to this case, provides evidence that Gen

[12] **Post**, for example, is uttered in contexts with very different questions at issue. In Context 1, 'what is it that people in Mail Room 6 do all day?' and, in Context 2, something like 'Who would sort mail from Antarctica?'

displays distinctive context-sensitivity.

While I don't share Sterken's intuitions regarding **A-Dobermans**, that example is not necessary for her argument. Sterken's claim is not that there are no contexts relative to which generics and their adverbially-quantified generalisation both shift their truth values. Adverbially-quantified generalisations can be context-sensitive for a variety of reasons (e.g, implicature, domain-restriction, and context-sensitivity of the subject or predicate). Her claim is that generics exhibit an additional context-sensitivity that is not shared by their adverbially-quantified counterparts. For this, it is sufficient that there are *some* contexts relative to which generics shift their truth values, while their adverbially-quantified counterparts do not. Indeed, Sterken presents other examples that I find more plausible. Take for example:

Frenchmen: Frenchmen eat horse meat.

This is plausibly true in a context that requires relatively few Frenchmen to eat horse meat, as when a contrast is set up with another nationality:

Context 1: Traditional French food differs from the traditional food of even their closest neighbours. Frenchmen eat horse meat, for example, whereas Englishmen find the idea incredible (my example).

Sterken cites another context in which the generic is intuitively false, rendering the negation true:

Context 2: A group of nutritionists is querying the unhealthy eating patterns of the French population ..."Frenchmen eat croissants and baguettes. They don't eat traditional food, like horse meat and grains" (Sterken, 2015a, pp. 314-5)

My ear agrees with Sterken here. The adverbial variants sound false in both contexts:

A-Frenchmen: Typically/Generally/Normally Frenchmen eat horse meat.

Frenchmen do not generally/typically/normally eat horse meat, even if it is a traditional food. While adverbially-quantified generalisations may exhibit some sensitivity to context, Sterken concludes that generics display an additional, distinctive context-sensitivity.

Sterken suggests that this distinctive generic context-sensitivity is best explained by analysing Gen as an indexical quantifier. The following section will present an alternative source of the intuitive difference in truth value between generics and their adverbially-quantified counterparts. It is worth bearing in mind that this alternative source of context-sensitivity is not intended to completely undermine the A-Quantifier Test. Indeed, I will make use of the test in section 5 to argue that the default generalisation account should explain our truth-value judgments through heuristics, rather than linguistic context-sensitivity. The test also provides some evidence for the indexical analysis, in that it successfully eliminates some competing accounts. It does not, however, eliminate all competitors and future applications of the test should bear in mind the full range of theories that can accommodate the data.

4 Towards an Alternative

4.1 The Default Mechanism of Generalisation

Leslie (2007, 2008) argues that humans have a prelinguistic mechanism of generalisation that associates kinds with properties.[13] With the acquisition of language, we learn alternative ways to generalise, such as those associated with the quantifiers 'all', 'most' and 'some'.[14] Leslie hypothesises that the prelinguis-

[13] See Graham, Kilbreath, & Welder (2004) for more on pre-linguistic generalisation.

[14] Hollander, Gelman, & Star (2002) suggest that this learning process takes place around four years of age. In their studies, three-year-olds responded to generics, existentials and generics in way that four-year-olds and adults respond only to generics.

tic mechanism remains our default mode of generalisation, however. Because generics include no explicit quantifier to override the default, they are interpreted through the prelinguistic mechanism.[15] Leslie therefore takes the default mechanism of generalisation to explain both our dispositions to utter generics and our dispositions to assess them as true:

> If a speaker's knowledge and experiences with members of a kind K leads her default mechanism to generalize the property of being F to that kind, then she will express this with the generic 'Ks are F'. Similarly, her hearer would judge the utterance to be true if, given his knowledge and experience, his default mechanism would generalize the property of being F to the Ks.
>
> Leslie (2008, p. 22)

If the default mechanism of generalisation is responsible for the shift in generics' apparent truth values, the stability of adverbially-quantified generalisations can be explained as a result of shifting away from the default mechanism.

Leslie argues that the default mechanism is sensitive to several contingent psychological factors. Leslie suggests, for example, that we are disposed to agree with generics only when the counterinstances are *negative*, and that we agree with generics more easily when they attribute *striking* or *characteristic* properties. These terms require some explanation.

First, counterinstances to 'Ks are F' are *negative* when Ks that are not F do not display any psychologically salient alternative property. So, for example, counterinstances to 'Lions have manes' are negative. The lions that lack manes don't have any particularly interesting alternative property: they simply lack manes. Counterexamples to 'Humans have dark hair', however, are positive. The people without dark hair exhibit salient alternatives properties, such as

[15] This view of generics as defaults is supported by Gelman and Brandone (2010), Gelman (2010), and Hollander, Gelman, & Raman (2009).

having red hair. According to Leslie, our default mechanism of generalisation only associates a property F with a kind K when the counterinstances to 'Ks are F' are negative. This explains why we are not disposed to agree with 'Humans have dark hair', despite the proportion of people with dark hair being far higher than the proportion of lions with manes.

Second, *characteristic properties*. The default mechanism is disposed to generalise about kinds more easily along certain dimensions. We expect, for example, that members of the same animal species share a characteristic diet, mode of reproduction, and so on. Where a generic concerns one of these characteristic properties, we generalise based on very little evidence. We might agree with 'Ducks lay eggs', for example, based on experience of a single egg-laying duck. We don't associate animal species with particular sexes, however, explaining why we are not disposed to agree with 'Ducks are female', despite there being more female ducks than egg-laying ducks.[16] Characteristic properties vary depending on the kind of kind in question. For artifacts, their characteristic property is their function. Leslie suggests that the default mechanism always associates a kind of artifact with its function, regardless of whether there are any instances that fulfil this function, explaining why we might agree with 'Cold fusion reactors produce power' even if one has never been constructed.

Finally, some properties are particularly *striking*. Leslie's usual examples are properties that pose a danger to humans. Where a property is particularly striking, Leslie suggests that our default mechanism associates it with a kind so long as some members of the kind have the property and all members of the kind are disposed to have it. This is intended to explain why we agree with 'Mosquitoes carry West Nile virus', despite the fact that very few mosquitoes are in fact carriers, but don't agree with 'Animals carry West Nile virus'. Leslie

[16] Our response to this case may in fact be overdetermined, given that male ducks constitute a positive counterinstance to 'Ducks are female' but a negative counterinstance to 'Ducks lay eggs'.

(2008, p. 41) assumes that all mosquitoes are disposed to carry West Nile virus but, of course, not all animals are.[17]

4.2 Heuristics for Generics

So, according to Leslie, we are disposed to agree with generics when our default mechanism associates the right property with the relevant kind and our default mechanism does so only under the following conditions:

Generic Truth: The counterinstances are negative,[18] 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, p. 43)

Leslie's view of generics so far involves three key claims:

1. There is a default mechanism of generalisation that associates kinds with properties.

2. The default mechanism associates kinds and properties under the conditions identified in **Generic Truth**.

[17] Leslie doesn't detail the required disposition but it faces significant problems. Note, for example, that we don't agree with 'Humans have HIV'. Perhaps some humans are immune to HIV (Ni, Wang, & Wang, 2018) but the required mutation is not known to most people. Indeed, the susceptibility of mosquitoes to viruses, and the degree to which they are disposed to transmit viruses, vary both between and within mosquito species (Hardy, 1988). See Sterken (2015b, pp. 2500-3) for reasons to doubt that Leslie's disposition requirement can be finessed to include only the intuitively true generics.

[18] Note that this can be read as a generic, albeit in definite rather than bare plural form. It isn't entirely clear, therefore, exactly what proportion of counterinstances must be negative, according to this condition. If some small number of lions had spikes in place of manes, would we judge 'Lions have manes' to be true?

3. When the default mechanism associates property F with kind K, we are disposed to treat the generic 'Ks are F' as true.

These claims characterise the heuristics associated with generics. The first posits a cognitive mechanism, the second describes the working of this mechanism, and the third identifies this mechanism as the way that we decide whether generics are true or false. If these heuristics are responsible for the context-sensitivity of our judgments about the truth of generics like **Frenchmen**, but adverbially-quantified sentences like those in **A-Frenchmen** are interpreted through different heuristics, then we have an explanation of our diverging truth-value judgments that doesn't posit semantic context-sensitivity.

To these three claims, Leslie adds a claim about the truth conditions of generics:

4. The truth conditions of generics are given by the conditions under which the default mechanism associates kinds and properties.

As Leslie (2008, p. 43) puts it, "Since this mechanism is responsible for our understanding of generics, providing an account of this mechanism has also allowed us to understand the circumstances in which generics are true or false." Sterken (2015b, p. 2494) describes Leslie's account as assuming that the "primitive cognitive mechanism of generalisation has certain accuracy conditions" and records in a footnote that a reviewer suggested "a better interpretation is that certain conditions make the primitive mechanism activate." Here, Sterken draws our attention to the application conditions of the theory, while the reviewer draws our attention to the heuristics.

In what follows, however, I will focus on Leslie's theory of heuristics, rather than her theory of truth conditions. Johnston and Leslie caution against drawing conclusions about application conditions from speakers' intuitions about correct usage. As they put it, "the method of appealing to our judgments as

to whether we should apply or withhold a term in a variety of imaginary cases is obviously a way of manifesting our criteria or ways of telling whether the term applies. It is not obviously a way of manifesting our 'implicit grasp' of the application conditions of terms" (Johnston and Leslie, 2012, pp. 135-6).[19] That is, intuitions about cases might be a good way of revealing the heuristics by which we judge whether a term applies but we cannot infer the application conditions of a term from these heuristics.

Leslie's inference from truth-value intuitions to truth-conditions is not obviously more acceptable than the inference that Johnston and Leslie (2012, pp. 134-6) identify as "a bad verificationist error": identifying the application conditions of 'dog' with the heuristics that we use to judge whether something is a dog. Those heuristics might work well enough in ordinary circumstances but fail when we encounter "A coiffed squirrel ... made to look like a chihuahua". Likewise, the heuristics that we use to judge whether a generic is true or false might work well enough in ordinary circumstances, but may also be prone to error, leading us to mistakenly treat false generics as true, or vice versa.

None of this is to argue that Leslie's truth conditional theory is incorrect. For the purposes of this paper, however, it is important to separate Leslie's theory of heuristics from her theory of truth conditions. The following section will set the truth conditional aspect of Leslie's theory to one side and consider whether Sterken's context-sensitivity data might be explained entirely by Leslie's heuristics.

5 An Alternative to Context-Sensitivity

Consider the following generic, uttered by a worker at Store A, which sells expensive coffee, many cheap second-hand books, and a few expensive rare

[19] See also Williamson (2020, p. 26).

editions, in response to a customer who asks what they can buy for two dollars:

Books: Books are cheap.

Suppose the worker goes into another store, Store B, on their day off. Store B sells cheap coffee and rare, highly-priced books. They ask what they can buy with limited funds and are told **Books**. 'That's not true!', our protagonist replies, and they leave the store in a state of incredulous shock.

How can Leslie's view explain the behaviour of our protagonist? The first utterance supposedly shows that our protagonist's default mechanism of generalisation associates the kind *books* with the property *being cheap*. If they associate this kind and property, however, they should be willing to accept the utterance of the assistant in Store B. It seems, therefore, that Leslie is committed to the view that our protagonist both does and does not associate books both with the property of being cheap. How can we make sense of this?

Suppose that 'Books' amounts to something like 'books in this store.' **Books** would then express different default generalisations in each context. In Store A, it would express an association between the kind *books in Store A* and the property *being cheap*. As our protagonist's default mechanism of generalisation associates this kind with the property, they would judge it true. In Store B, the generic would express an association between the kind *books in Store B* and the property *being cheap*. As our protagonist's default mechanism of generalisation does not associate this different kind with the property, they would judge it false. Alternatively, the predicate could be context-sensitive, amounting to something like 'cheap in this store'. Again, **Books** would then express different default generalisations in different contexts. In Store A, it would express an association between the kind *books* and the property *being cheap in Store A*. As our protagonist's default mechanism of generalisation associates the kind with this property, they would judge it true. In Store B, the generic expresses

an association between the kind *books* and the property *being cheap in Store B*. As our protagonist's default mechanism of generalisation does not associate the kind with this different property, they would judge it false. Either way, we would have an explanation of why our protagonist is willing to utter **Books** to their customers in Store A, but not willing to accept its utterance by the assistant in Store B.

According to Sterken's A-Quantifier Test, if the subject or the predicate is responsible for **Books** changing its apparent truth value between contexts, then we should expect the same difference to hold for its adverbially-quantified variants:

A-Books: Typically/Generally/Normally books are cheap.

It seems to me that these adverbially-quantified generalisations do shift their truth value between contexts. In Store A, books are typically cheap and **A-Books** is true. In Store B, books are not typically cheap and **A-Books** is false. So we can explain why **Books** and **A-Books** change their apparent truth values across contexts by positing context-sensitivity in either the subject or predicate. The same explanation cannot easily be extended to **Frenchmen**, however. If context-sensitivity of either the subject or predicate were responsible for **Frenchmen** changing its apparent truth value between contexts, we would also expect **A-Frenchmen** to change its apparent truth value across contexts but it does not.

An alternative explanation is available, however: the kinds and properties associated by the default mechanism of generalisation change between contexts. As applied to **Books**, when in Store A, the worker associates the kind *books* with the property *being cheap* but this changes when they move to Store B. Changing their associations allows them to navigate two different environments. The worker follows exactly the same heuristic in both contexts, relying on their

default mechanism of generalisation, but their truth-value judgments change as a function of changes in their associations. Likewise, in Context 1, the kind *Frenchmen* is associated with the property of *eating horse meat*. If the associations of the default mechanism vary between contexts, we can explain different truth-value judgments in different contexts without positing any semantic context-sensitivity. In both contexts, the subject picks out the same kind and the predicate picks out the same property but that property is cognitively associated with that kind in some contexts and not in others. This leads to differing truth-value judgments because our heuristics for assessing truth and falsity operate based on these contextual associations.

Leslie's account already has some flexibility with respect to the conditions under which kinds and properties are associated.[20] **Generic Truth** tells us, for example, that when a property F lies along a characteristic dimension for some kind K, the property is associated with the kind so long as some Ks are F. It isn't clear, however, that these characteristic dimensions are always independent of context. Context 1 sets up a context in which we associate nationalities with traditional cuisine. In that context, traditional cuisine is thought of as a characteristic dimension and the kind *Frenchmen* is associated with the property *eating horse meat*. In the second context, however, traditional cuisine is explicitly rejected as a basis for generalisation and isn't considered a characteristic dimension. Instead, associations are made based on the prevalent eating habits of the population.

The view that generics are affected by context has the resources to account for the results of the A-Quantifier Test. To account for the results, we need two things. First, we need to explain how judgments about truth-values of generics could vary between contexts. Second, we need to explain how judgments about the truth-values of their adverbial variants could remain stable across the same

[20]Thanks to an anonymous reviewer for this suggestion.

contexts. The first is secured by the possibility of contextual variation in associations. The second is secured by the notion of a default mechanism. According to Leslie's theory, the very purpose of quantifiers is to shift interpretation away from the default mechanism of generalisation.[21] When assessing 'Generally, Frenchmen eat horse meat', for example, the term 'generally' functions to shift interpretation from the prelinguistic default mechanism to whatever mode of generalisation we have learned to associate with 'generally'. Even if I currently associate the property *eating horse meat* with the kind *Frenchmen*, I can judge 'Generally, Frenchmen eat horse meat' to be false because I take 'generally' to require a majority of cases to conform to the generalisation and I do not believe that most Frenchmen eat horse meat.[22]

The hypothesis that associations vary between context is not unrealistic. Contextual effects have been found for many conceptual processing tasks, including word recall (Barclay et al., 1974; Anderson & Ortony, 1975; Greenspan, 1986; Zeelenberg, 2005), object identification (Biederman, Mezzanotte, & Rabinowitz, 1982; Palmer, 1975; Boyce, Pollatsek, & Rayner, 1989; Murphy and Wisniewski, 1989), property generation (Tabossi and Johnson-Laird, 1980; Barsalou, 1993, p.32), typicality judgments (Barsalou, 1993, p. 33; Roth and Shoben, 1983), lexical decision (Greenspan, 1986; Kellas et al., 1991), reading times (Tabossi and Johnson-Laird, 1980), association judgements (Greenspan, 1986), property verification (Barsalou, 1982), generic agreement (Tabossi and Johnson-Laird,

[21] That is not to say, however, that this function is always fulfilled. Leslie, Khemlani, & Glucksberg (2012), for example, observed the "generic overgeneralization effect" in which English speakers agreed to seemingly false sentences like 'All ducks lay eggs'. They take this as evidence in favour of the default generalisation hypothesis. If quantifiers aim to override the generic default, then we should expect failures in which speakers nevertheless interpret quantified statements as generics.

[22] I actually take the generalisation to be slightly more complicated. What is it, after all, for an individual to eat horse meat and so count as conforming to the generalisation? Is it for them to have eaten horse meat once in their lives? Regularly? How regularly? Is a bite on special occasions, barring exceptional circumstances, sufficient? The inclusion of 'generally' does not eliminate all elements of genericity from the sentence. My suspicion is that I judge the sentence false because I assume there are sufficiently many Frenchmen who I would not characterise as generic horse-meat-eaters.

1980) and category judgments (Tobia, Newman, & Knobe, 2020; Machery & Seppälä, 2011; Zarl and Fum, 2014).[23]

Machery & Seppälä (2011), for example, found that many subjects were willing to assent to seemingly incompatible category judgments e.g., 'tomatoes are vegetables' and 'tomatoes are not vegetables' when both were qualified with 'In a sense'. This may be because they are aware that there are some contexts (e.g., biology) in which they would not associate the kind *tomatoes* with the property of *being a vegetable* but there are other contexts (e.g., a grocery store) in which they would. This interpretation is strengthened by experiments conducted by Zarl and Fum (2014), who found that participants are more willing to accept both statements when they are presented individually, rather than in the same context, and that more explicit sentential context (e.g., 'In a technical sense...') reduced the number of incompatible judgments.

The hypothesis that contextually-varying cognitive associations determine our contextually-varying truth-value judgments about generics is speculative, requiring empirical support. Regardless, our default position should be that contextual differences in truth-value judgments about generics result from *some* effect of context on heuristics, not that generics are semantically context-sensitive. Competent speakers can disagree about the truth-values of generics. As noted in section 3.3, for example, Sterken and I disagree about the truth-value of **Dobermans** in some contexts. Parties to such a disagreement cannot both be right. Given that competent speakers can so easily be wrong, they must be relying on fallible heuristics, rather than, say, semantic knowledge encoded within their brains.[24] If that is the case, then differences in truth-value judgments must result from contextual effects on those heuristics and it is a further claim, re-

[23] See the literature review in (Yeh and Barsalou, 2006) for further examples of these effects.

[24] We can get things wrong, even when they are encoded within our brains. We might fail to retrieve this information accurately, for example, due to inattentiveness or high cognitive load. Given that these disagreements can arise for competent speakers in ideal conditions, however, I claim that only the use of heuristics could explain them.

quiring further argument, that these heuristics identify the actual truth-values of the generics and that generics are therefore context-sensitive. Indeed, any substantial context-sensitivity argument will be about expressions for which disagreements can arise, suggesting that our truth-value judgments must be the result of fallible heuristics. The burden of proof is therefore on the author of any such argument to show that these heuristics deliver the correct result in the cases central to their argument.

6 Conclusion

Arguments for context-sensitivity usually begin with data about our judgments of truth and falsity in different contexts. Often, this data is suppressed in that authors talk directly about the truth values, but these assertions must be rooted in their *judgments* about truth values. This paper has presented a way of explaining these judgments without positing context-sensitivity. Judgments about the truth and falsity of sentences depend directly on the heuristics that we use to judge truth and falsity. These heuristics may lead to different truth-value judgments in different contexts.

This explanation has been illustrated through Sterken's argument for her indexical analysis of generics. Leslie has suggested that generics might be interpreted through heuristics that differ from those we use to interpret adverbially-quantified generalisations. Contextual variations might therefore affect our the former but not the latter. While this is not intended to completely undermine Sterken's A-Quantifier Test, it does limit the degree to which that test provides support for Sterken's indexical analysis of generics.

My hope is that future authors will pay more attention to the role of heuristics in future discussions of context-sensitivity. The key point is this: Intuitions about truth and falsity are evidence most directly for the heuristics that we use

to judge truth and falsity. Where some contextual difference affects the output of these heuristics, our truth-value judgments can change without any corresponding change in truth conditions. In future, authors of context-sensitivity arguments should consider this possibility and either provide reason to think that the relevant truth-value judgments are not the result of fallbile heuristics, or reason to think that these heuristics are accurate in the cases central to their argument.

References

Anderson, R. C., & Ortony, A. (1975). On putting apples into bottles — A problem of polysemy. *Cognitive Psychology* 7(2), 167–180.

Barclay, J. R., Bransford, J. D., Franks, J. J., McCarrell, N. S., & Nitsch, K. (1974). Comprehension and semantic flexibility. *Journal of Verbal Learning & Verbal Behavior* 13(4), 471–481.

Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition* 10(1), 82–93.

Barsalou, L. W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. F. Collins, S. E. Gathercole, M. A. Conway, & P. E. Morris (Eds.), *Theories of memory* (pp. 29–101). Lawrence Erlbaum Associates, Inc.

Biederman, I., Mezzanotte, R. J., & Rabinowitz, J. C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology* 14(2), 143–177.

Boyce, S. J., Pollatsek, A., & Rayner, K. (1989). Effect of background infor-

mation on object identification. *Journal of Experimental Psychology. Human Perception and Performance* 15(3), 556–566.

Cohen, A. (2004). Existential Generics. *Linguistics and Philosophy* 27(2), 137-168.

Fennell, J. (2013). "The Meaning of 'Meaning is Normative' " *Philosophical Investigations* 36, 56-78.

Gabrielcik, A. and Fazio, R.H. 1984. Priming and Frequency Estimation: A Strict Test of the Availability Heuristic. *Personality and Social Psychology Bulletin* 10(1), 85-89.

Gelman, S.A. and Brandone, A. C. (2010). Fast-Mapping Placeholders: Using Words to Talk About Kinds. *Language Learning and Development* 6(3), 223-240.

Gelman, S.A. (2010). Generics as a Window onto Young Children's Concepts. In F.J. Pelletier (Ed.), *Kinds, Things, and Stuff: Mass Terms and Generics* (pp. 100-120). Oxford: Oxford University Press.

Graham S.A, Kilbreath C.S, & Welder A.N. (2004). Thirteen-Month-Olds Rely on Shared Labels and Shape Similarity for Inductive Inferences. *Child Development* 75(2), 409-27.

Greenberg,Y. (2007). Exceptions to Generics: Where Vagueness, Context Dependence and Modality Interact. *Journal of Semantics* 24(2), 131-167.

Greenspan, S. L. (1986). Semantic flexibility and referential specificity of concrete nouns. *Journal of Memory and Language* 25(5), 539–557.

Grice, P. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

Hardy, J. L. (1988). Susceptibility and resistance of vector mosquitoes. In T.P., Monath (Ed.), *Arboviruses: epidemiology and ecology, vol. I* (pp. 87–126). Boca Raton, FL: CRC Press.

Hattiangadi, A. (2006). Is Meaning Normative? *Mind & Language* 21, 220–240.

Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. Doctoral dissertation. University of Massachusetts at Amherst.

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

Hollander, M.A, Gelman, S.A., & Raman, L. (2009). Generic Language and Judgements About Category Membership: Can Generics Highlight Properties as Central? *Language and Cognitive Processes* 24(4), 481–505.

Johnston, M. and Leslie, S.J. (2012). Concepts, Analysis, Generics, and the Canberra Plan. *Philosophical Perspectives* 26, 113-171.

Johnston, M., & Leslie, S.J. (2019). Cognitive Psychology and the Metaphysics of Meaning. In A.I. Goldman and B.P. McLaughlin (Eds.), *Metaphysics and Cognitive Science* (pp. 183-205). Oxford: Oxford University Press.

Kahneman, D. (2012). *Thinking Fast and Slow*. London: Penguin Books.

Kamp, H. (1981/1984). A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, & M. Stokhof (Eds.), *Truth,interpretation, and information* (pp. 1–41). Dordrecht: Foris.

Kellas, G., Paul, S. T., Martin, M., & Simpson, G. B. (1991). Chapter 3 Contextual Feature Activation and Meaning Access. In G. B. Simpson (Ed.), *Advances in Psychology* (Vol. 77, pp. 47–71). North-Holland.

Knobe, J. (2003). Intentional action in folk psychology: An experimental investigation. *Philosophical Psychology* 16(2), 309-324.

Kripke, S.A. (1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press.

Krifka, M., Pelletier, F.J., Carlson, G., ter Meulen, A. Chierchia, G. & Link, G. (1995). Genericity: An Introduction. In G.N. Carlson & F.J Pelletier (Eds.), *The Generic Book* (pp. 1-124). University of Chicago Press. .

Leslie, S.J. (2007). Generics and the Structure of the Mind. *Philosophical Perspectives* 21(1), 375-403.

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

Leslie, S.J., Khemlani, S., & Glucksberg, S. (2011). Do All Ducks Lay Eggs? The Generic Overgeneralization Effect. *Journal of Memory and Language* 65(1), 15-31

Lewis, D. (1975). Adverbs of Quantification. In E. L. Keenan (Ed.), *Formal Semantics of Natural Language* (pp. 3-15). Cambridge: Cambridge University Press.

Machery, E., & Seppälä, S. (2011). Against Hybrid Theories of Concepts. *Anthropology & Philosophy* 10, 99–127.

Murphy, G. L. and Wisniewski, E. J. (1989). Categorizing objects in isolation and in scenes: What a superordinate is good for. *Journal of Experimental Psychology. Learning, Memory, and Cognition* 15(4), 572–586.

Ni, J., Wang, D. & Wang, S. (2018). The CCR5-Delta32 Genetic Polymorphism and HIV-1 Infection Susceptibility: a Meta-analysis. *Open Med (Wars)* 13, 467–474.

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

Nickel, B. (2016). *Between Logic and the World: An Integrated Theory of Generics*. Oxford: Oxford University Press.

Palmer, S. E. (1975). The effects of contextual scenes on the identification of objects. *Memory & Cognition* 3(5), 519–526.

Putnam, H. (1975). The meaning of 'meaning'. *Minnesota Studies in the Philosophy of Science* 7, 131-193.

Rosch, E.H. (1973). Natural categories. *Cognitive Psychology* 4(3), 328-350.

Rosch, E. (1978). Principles of Categorization. In E. Rosch & B.B. Lloyd (Eds.) *Cognition and Categorization* (pp. 27–48). Hillsdale: Lawrence Erlbaum.

Rosch, E. & Mervis, C.B. (1975). Family Resemblances: Studies in the Internal Structure of Categories. *Cognitive Psychology* 7(4), 573–605.

Roth, E. M. and Shoben, E. J. (1983). The effect of context on the structure of categories. *Cognitive Psychology* 15(3), 346–378.

Sterken, R. K. (2015a). Generics in Context. *Philosophers' Imprint*, 15.

Sterken, R. K. (2015b). Leslie on Generics. *Philosophical Studies* 172(9), 2493-2512.

Tabossi, P. and Johnson-Laird, P. N. (1980). Linguistic Context and the Priming of Semantic Information. *Quarterly Journal of Experimental Psychology* 32(4), 595–603.

Tobia, K. P., Newman, G. E., & Knobe, J. (2020). Water is and is not H_2O. *Mind & Language* 35(2), 183–208.

Tversky, A. and Kahneman, D. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185(4157), 1124-1131.

Wikforss, Å.M. (2001). Semantic Normativity. *Philosophical Studies* 102, 203–226

Williamson, T. (2020). *Suppose and Tell: The Semantics and Heuristics of Conditionals.* Oxford: Oxford University Press.

Yeh, W. and Barsalou, L.W (2006). The Situated Nature of Concepts. *The American Journal of Psychology* 119(3), 349-384.

Zarl, F., & Fum, D. (2014). Theories of Concepts and Contradiction Acceptance. In P. Bernardis, C. Fantoni, & W. Gerbino (Eds.), *TSPC2014: Proceedings of the Trieste Symposium on Perception and Cognition, November 27th-28th 2014* (pp. 157–161). Trieste, Italy: EUT Edizioni Università di Trieste.

Zeelenberg, R. (2005). Encoding specificity manipulations do affect retrieval from memory. *Acta Psychologica* 119(1), 107–121.