Generics, Content and Cognitive Bias

Rachel Katharine Sterken*
University of Oslo

January 15, 2015

Paradigm generics communicate quite strong generalisations:

(1) Dogs bark.

(2) Tigers have stripes.

(3) Candy is bad for your teeth.

Indeed, it is only a small minority of mute dogs that don't bark, all tigers have stripes unless they are born with some sort of birth defect, and most candy contains sugar or sweeteners and thus, is detrimental to dental health.

Despite the paradigm, there are numerous generics that appear to express quite weak generalisations — consider the following:

(4) Mosquitoes carry the West Nile virus.

(5) Sharks attack bathers.

(6) Rottweilers maul children.

(7) Tigers eat people.

(8) Australian snakes are poisonous.

Examples like (4)-(8) are intuitively true even though only a minority, even a very small minority, of the kind satisfies the predicated property: (4) is intuitively true even though less than one percent of mosquitoes in fact carry the West Nile virus. (5), (6) and (7) are

*Acknowledgments: This paper has benefited from discussions with Nicholas Allott, Mahrad Almotahari, Nicholas Asher, Herman Cappelen, Gennaro Chierchia, Ephraim Glick, John Hawthorne, Torfinn Huvenes, Philipp Koralus, Nick Kroll, Sarah-Jane Leslie, David Liebesman, Michael Morreau, Bernhard Nickel, Jeff Pelletier, Jennifer Saul, Jonathan Schaffer, Martin Smith, Andreas Stokke, Brian Weatherson and Elia Zardini. Parts of the paper were presented at a Workshop on Philosophy of Language at Harvard University, the Society for Exact Philosophy (SEP) Conference held at the University of Alberta, the Harvard/MIT Graduate Conference, the Yale/UConn Graduate Conference, the Princeton/Rutgers Graduate Conference, the University of Oslo (CSMN) and the University of St Andrews (Arché). I thank the audience members and organisers at these events. I also thank an anonymous referee for this journal for helpful comments.

intuitively true even though these animals are only violent when put in the appropriate, quite rare circumstances — e.g., if a child is taunting or bothering a Rottweiler, then the Rottweiler might react by mauling the child. Further, there is an intuitively true reading of (8), even though only about a dozen out of the 170 species of snake found in Australia are venomous enough to be dangerous to humans.

Examples like (4)-(8) have made the task of theorising about the truth-conditions of generics especially difficult: It is difficult to simultaneously account for the strength of (1)-(3) and the weakness of (4)-(8).

Theorists have made numerous proposals about how to deal with cases such as (4)-(8).[1] In what follows, I motivate and defend an error theory to deal with such cases. In other words, I argue that our intuitions about (4)-(8) are mistaken, and that a theory of generics should not predict that they are true. In section 1, I briefly outline the existing proposals for how to deal with (4)-(8) and say what I take to be problematic about them. In section 2, I outline my own proposal and the linguistic evidence for it. In section 3, I use this class of examples as a case study to draw out some more general conclusions about when error theories are needed in semantics.

1 Two Proposals

1.1 Capacity Readings:

One strategy for dealing with (4)-(8) is to treat those sentences as ambiguous. On their generic interpretation, (4)-(8) are false, but on their other, their preferred, interpretation, they are true. The most prominent such proposal is that their preferred interpretation is a generic *capacity reading*. This proposal is defended in different forms in Asher and Pelletier (2012), and Nickel (ms.b, ms.a).[2] According to this proposal, the preferred interpretation of (4), for example, is akin to one of the following:

(9) a. All normal mosquitoes can carry the West Nile virus.
b. All normal mosquitoes have the capacity to carry the West Nile virus.

(10) a. Generally, mosquitoes can carry the West Nile virus.
b. Generally, mosquitoes have the capacity to carry the West Nile virus.

The generic capacity reading is weaker than the standard generic interpretation because it contains the capacity modal, whereas the standard generic interpretation does not.

[1] See, e.g., Carlson (1977), A. Cohen (2001), Leslie (2007a, 2008), Liebesman (2011), Asher and Pelletier (2012) and Nickel (ms.a, ms.b).

[2] Another proposal is that their other interpretation is a *relative reading* — e.g., the relative reading of (4) is akin to *mosquitoes are more likely than other types of insect to carry the West Nile virus.* This proposal is defended in A. Cohen (1996, 2001). For compelling arguments against this proposal, see Leslie (2007b).

A further proposal often advocated by my informants is that (4)-(8) receive what is called a *reverse interpretation.* On such a proposal, (4) is akin to *all normal West Nile carriers are mosquitoes* or *generally West Nile carriers are mosquitoes.* For compelling argument against such a proposal see A. Cohen (2001).

There is independent motivation for the existence of capacity readings, and so it seems like an easy enough fix.[3] However, there are at least three objections to this proposal.

The first trouble is that capacity readings are normally only available with *eventive* predicates — which indeed, the predicates in at least (4)-(7) are. Capacity readings are not normally available with non-eventive predicates. For example, it is hard to find contexts in which the preferred interpretation of (2) is *tigers can have stripes* or *tigers have the capacity to have stripes*. A problem with the capacity reading proposal, then, is that there are examples like (4)-(7) which contain non-eventive predicates — consider again (8) and the following:

(13) Mosquitoes are carriers of the West Nile virus.

Both (8) and (13) are intuitively true, weak generics and yet, without some substantive additional theorising, they cannot be construed as generic capacity readings. With the absence of an event variable in the predicates in (8) and (13), the desired reading is not available. In other words, (8) cannot mean that all snakes normally can be poisonous or that generally snakes have the capacity to be poisonous, and similarly, (13) cannot mean that all normal mosquitoes can be carriers of the West Nile virus or that generally mosquitoes have the capacity to be carriers of the West Nile virus.[4]

If you find yourself unconvinced by examples (8) and (13), or you find that this is equally an argument against the claim that capacity readings only occur with eventive predicates, then I offer the following two tests as a means to distinguish the capacity reading from the generic reading. If (4)-(8) are generic ascriptions of capacities, then they should pass my proposed tests. Since they do not, I conclude that (4)-(8) are regular generics after all.[5]

Test 1: Actuality Entailments and Capacities

One feature of explicit capacity and ability attributions, like (14a) and (15a) below, is that they do not generally entail that the event related to the embedded predicate took place, will take place or is taking place.[6] Consider the following:

[3] Nickel (ms.b) provides the following illustrative example:

(11) The Eurostar goes 120 mph.

(12) My Peugeot goes 120 mph.

The pair of examples is meant to illustrate that the predicate *goes 120mph* is ambiguous between a habitual and capacity reading: Intuitively, (11) says what speed the Eurostar normally goes and (12) says what speed the speaker's Peugeot can go. Thus, there is a class of predicates which are ambiguous between a habitual interpretation and capacity interpretation.

[4] Of course, one could argue for a different understanding of the availability of the capacity reading — e.g., an unpronounced modal. I will not discuss such options: I hope that the subsequent discussion will forego the need to consider more detailed options.

[5] Again, my aim here is not to argue that generic capacity readings aren't ever available, rather my aim is to argue that the preferred interpretations of (4)-(8) are not generic capacity readings.

[6] See, e.g., Bhatt (1999).

(14) a. Yesterday, it could have rained here. / Yesterday, it was able to rain here.
b. Does not entail: Yesterday, it rained here.
c. It could have rained yesterday, but it didn't.

(15) a. The mailman can deliver the mail. / The mailman is able to deliver the mail.
b. Does not entail: The mailman delivered/will deliver/is delivering the mail.
c. The mailman can deliver the mail, but he didn't/he won't/he isn't.

The capacity and ability attributions in (14a) do not entail (14b), nor does the conjunction in (14c) sound contradictory or infelicitous in some way. Similarly, in the case of (15). Given this property of capacity and ability attributions, the following acts as a test: If a conjunction with (4)-(8) as conjuncts analogous to those in (14c) and (15c) sound bad, then this is evidence that (4)-(8) are not capacity readings. Applying the test, consider:

(16) a. Mosquitoes carry the West Nile virus. # But they don't because they've been irradiated.
b. Mosquitoes can carry the West Nile virus. But they don't because they've been irradiated.

(17) a. Sharks attack bathers. # But they don't because there are shark nets on all beaches.
b. Sharks can attack bathers. But they don't since there are shark nets on all beaches.

The conjunctions in (16b) and (17b) where one of the conjuncts is an explicit capacity reading sound fine. By contrast the conjunctions in (16a) and (17a) where one of the conjuncts is simply a generic (i.e., (4) and (5) respectively) sound contradictory or at least infelicitous in some way. Thus, given our test, this is evidence that cases like (4)-(8) are not capacity readings.

Test 2: Conditional Restriction

In recent work, Kratzer (2012) claims that *if*-clauses cannot restrict capacity and ability modals. She provides the following illustrative examples:[7]

(18) If he has a kitchen, he can cook.

(19) If I was taller, I could reach the ceiling.

(18) is not naturally read as saying of someone that he has the ability to cook on the condition that he has a kitchen. Similarly, (19) is not naturally read as saying that I have the ability to reach the ceiling conditional on being taller. These would be quite bizarre capacities. Given examples like (18) and (19), Kratzer concludes that *if*-clauses do not restrict capacity and ability modals. But, of course, *if*-clauses can restrict plain generics — consider:

(20) If shops are open on Sundays, then they get customers on Sundays.

[7] See Kratzer (2012, p.107).

(21) If cats jump, then they land on their feet.

(20) is naturally read as saying that shops get customers on Sundays given that they are open on Sundays, and (21) as saying that cats land on their feet conditional on them jumping.

Thus, if Kratzer is correct,[8] then we can check whether or not (4)-(8) are capacity readings by using the following test: If *if*-clauses can restrict (4)-(8), then (4)-(8) are not capacity readings, rather (4)-(8) are plain generics. To this end, consider the following:

(22) If mosquitoes are found near birds, they carry the West Nile virus.

(23) If sharks aren't given adequate space, they attack bathers.

(22) is naturally read as saying that mosquitoes carry the West Nile virus conditional on being found near birds, and (23) as saying that sharks attack bathers conditional on them not having enough space. Hence, *if*-clauses can restrict cases like (4)-(8) and our test reveals that generics like (4)-(8) are not capacity readings.

I've presented three reasons to think that the preferred interpretation of (4)-(8) is not the capacity reading as Nickel, and Asher and Pelletier contend. Next, I consider a second prominent proposal for how to accommodate (4)-(8) in a semantics for generics, due to Leslie.[9,10]

Leslie on Mosquitoes and Sharks:

Leslie (2007a, 2008) observes that there seems to be something about (4)-(8) which distinguishes them from paradigm generics. In each of (4)-(8), the predicated property is somehow a *striking*, *alarming* or *dangerous* property for the kind to possess:

> The examples... have something in common: In all of them, the sentence attributes harmful, dangerous, or appalling properties to the kind. More generally, if the property in question is the sort of property of which one would be well served to be forewarned, even if there were only a small chance of encountering it, then generic attributions of the property are intuitively true. (2008, p.15)

Indeed, (4)-(8) do seem to pattern along these lines — it's not hard to generate more examples which follow suit:

[8] Schaffer and Szabo (2014) claim in a footnote that Kratzer's conclusion is too strong. They claim that ability modals can be restricted by *if*-clauses and that her data warrants a pragmatic explanation.

[9] There is consensus that domain restriction or tinkering with the modal properties of *Gen* are not viable solutions to deal with these cases. See, e.g., Leslie (2007a) and Asher and Pelletier (2012).

[10] It is worth noting that the (4)-(8) data is unproblematic for theories which treat generics as kind-predications. Indeed, the (4)-(8) data is treated as evidence that generics are a species of kind-predication. I will not discuss this option as it is beyond the scope of this paper to argue against a kind-predication analysis. I and others do that elsewhere. See, e.g., Carlson (1989), Leslie (2013) and Sterken (ms.).

(24) Earthquakes kill people and cause enormous damage.

(25) Airplanes crash.

(26) Smoking kills.

There are very few earthquakes that actually result in fatalities and large scale damage, and yet we judge (24) as true. Those that are afraid of flying often cite (25) as justification for their fears even though very few airplanes crash. (26) is asserted on many anti-smoking campaigns, though the percentage of smokers that die from smoking related causes each year is roughly 0.6-0.7 percent. Thus, Leslie's observation does seem to track the class of examples at issue. For the sake of brevity, I will follow Leslie and label such properties, *striking properties*.

To support her observation, Leslie cites several studies in empirical psychology which demonstrate a connection between our judgments of frequency and the strikingness of the information relevant to making those judgments. Rothbart, Fulero, Jenson, Howard, and Birrell (1978), for instance, tested the effects of striking information on estimates of statistical frequencies. The subjects in the study were presented with descriptions of fifty people, forty of the people were described as well-behaved, and ten of the people were described as engaged in striking negative criminal behaviour. The subjects were split into two groups: The subjects in the first group were given descriptions where the negative criminal behaviour of the ten people was especially appalling — like rape; while the subjects in the second group were given descriptions where the negative criminal behaviour of the ten people where by comparison much more mild — like shoplifting. The subjects were then asked to estimate the frequency of criminals in the group of fifty people they had just had described to them. The first group of subjects, who heard descriptions of the especially appalling criminal behaviour, provided higher estimates of the number of criminals in the group of fifty people described. The estimates of the first group of subjects were, thus, influenced by the strikingness of the information provided.

In another study, Cimpian, Brandone, and Gelman (2010) found that generic generalisations involving properties that were described as distinctive or dangerous were accepted more often than generic generalisations of other similar properties.

A further study, conducted by Abelson and Kanouse (1966), found that some generics, like (4)-(8), require very little evidence for acceptance, and yet "once accepted psychologically they appear[ed] to be commonly taken in a rather strong sense, as though the quantifier *always* had implicitly crept into their interpretation".

This provides strong evidence that there is at least some connection between our intuitions that (4)-(8) are true and the strikingness of the predicated properties at issue.[11]

Though I think Leslie's observation about the connection between the (4)-(8) cases and strikingness is apt — I won't take issue with the results empirical psychology, but merely the significance of the results — her account of the truth-conditions of such cases is not correct. Leslie proposes that we take our intuitions about such cases at face value

[11] It is worth noting that the capacity reading proposal does not explain this connection.

and incorporate them into our account of the truth-conditions of generics. She claims that when the property in question is striking, the truth-conditions of the generic are weak — in the sense that it suffices that merely some members of the kind possess the striking property for the generic to come out as true. In addition, it is a condition on their truth, that the counter-instances, those members of the kind that do not possess the predicated property, are at least disposed to have the striking property and that they do not have any positive alternative property.[12] In sum, Leslie provides the following truth-conditions for generics like (4)-(8):

> *Ks are F* is true if:
>
> > (i) the counter-instances (if any) are negative and;
> >
> > (ii) if F is striking, then some Ks are F and the others are (typically) disposed to be F.

However, this cannot be right. In the first place, it is a rather surprising position to claim that our theory of generics should take as evidence, intuitions which are subject to empirically demonstrated bias. That aside, the most pressing problem for Leslie's account is that there are counterexamples to the proposed sufficiency conditions. Consider, for instance, the following generics which satisfy the conditions outlined above and yet we judge them as false in normal contexts:

(27) Humans kill themselves.

(28) Homeowners start fires at night.

I take it when (27) and (28) are uttered on their own, they are not intuitively true.[13] Yet killing themselves and starting fires at night are quite striking properties for human beings and homeowners to possess, respectively. Moreover, the counter-instances — i.e., the humans and homeowners who don't satisfy the predicated property — are appropriately negative. That is, there is no perceptually salient or striking alternative property which the non-suicidal humans and non-pyro homeowners satisfy. Thus, it seems that Leslie's conditions are satisfied, making (27) and (28) true even though they are intuitively false.[14]

How can we account for Leslie's observation that strikingness plays an important role, if not directly in the semantics of generics? In the next section, I argue that an error theory is the best way to do this.

---

[12] Leslie introduces a distinction between positive and negative alternatives to a property F. An alternative property to F is positive if it is as equally salient, memorable and striking as F.

[13] Given the right context, they can come out as intuitively true — for example, consider:

(29) A: What surprising things do humans do?
B: Humans kill themselves.

However, this observation doesn't help Leslie. The relevant property should be striking in both contexts.

[14] For a more elaborate argument that such cases count as counterexamples, see Sterken (forthcoming). The central purpose of this paper is to present an alternative account and so the focus is not on refuting the alternatives.

2 An Error Theory

I propose that intuitions about the truth of cases like (4)-(8) should not be taken at face value. Their interpretation is defective: When we intuit (4)-(8) as true, we are making a mistake. It should not be a condition on a theory of genericity that it makes generics like (4)-(8) come out as true. We might have intuitions to the effect that they are, but these intuitions, I claim, are mistaken.

Having said that though, we cannot simply deny the intuitions that such generics are true, we do have them after all. What accounts for the fact that we are making errors? There are obviously many ways of making mistakes. I will focus on how cognitive biases, like our sensitivity to strikingness, can cause us to make errors. It could be that such cognitive biases interfere with how we form beliefs about the world, causing us to have incorrect beliefs — in this case, speakers exhibit a certain *blindness or ignorance about the world*. Alternatively, it could be that such factors interfere with how speakers interpret generics, causing us to get the content of the generic utterance wrong — in this case, speakers exhibit a *semantic* or *interpretive* error. In the former, the speaker mistakenly thinks that the generic is true because he thinks that in general mosquitoes carry the West Nile. In the latter case, the speaker mistakenly takes the content of the generic utterance to be weaker than it should be — e.g., taking it to have existential force so that the speaker thinks his utterance is true even though what it says is false. For simplicity, I will focus primarily on the former, however, I am officially neutral on whether or not the error is worldly or semantic, it could indeed vary between circumstances and cases.[15]

An error theory, thus, serves as a vindication of the connection between cases like (4)-(8) and strikingness. The empirical studies cited above are consistent with an error theory: The studies do not say what the truth-conditions of generic sentences are, rather they demonstrate a bias in how the mind processes information when it is striking to us — this demonstrated bias can be treated as impacting the truth-conditions of generic sentences (Leslie's view) or as merely impacting our judgments about the truth of such sentences.

Moreover, if a speaker thinks that, for example, the generic (4) is true and goes about her thinking as if it has a strong interpretation, then when thinking about mosquitoes, her tendency to overgeneralise will extend to sentences with explicit quantifiers. That is, she will also often be inclined to accept *many mosquitoes carry the West Nile* or *lots of mosquitoes carry the West Nile* as true, and *mosquitoes don't carry the West Nile* as false. This is what one would expect on an error theory.

At the same time, an error theory does not force any drastic claims about the truth-

[15] Some readers have found the distinction between errors about the world and interpretative errors confusing. The general idea is simply that the error could go two ways: in the former case (the error-about-the-world case) it does not affect the interpretation of the generic statement, just the agent's judgment about the world. In the latter case, the agent is lead to misinterpret the statement and, given the misinterpreted content, the statement is judged true (the agent is right about the world, but not about the content of the uttered sentence). The focus in what follows is on the former case: so they are all treated as generic statements and the error concerns the what the world is like.

conditions of generics or the nature of genericity. Not only that, but an error theory allows these troublesome cases to be set aside in theorising about generics.

Next, I present two types of linguistic evidence and an additional argument involving a variant of a test proposed by Kripke, all of which support an error theory.

2.1 Evidence 1: Contradictory Conjunctions

Consider the following sentences:

(30) a. Mosquitoes carry the West Nile virus, but typically they don't.
b. Typically mosquitoes don't carry the West Nile virus, but mosquitoes carry the West Nile virus.

(31) a. Sharks attack bathers, but generally they don't.
b. Generally sharks don't attack bathers, but sharks attack bathers.

(32) a. Rottweilers maul children, but normally they don't.
b. Normally Rottweilers don't maul children, but Rottweilers maul children.

(30)-(32) sound contradictory or at least perplexing in some way, but according to those advocating competing proposals, (30)-(32) could turn out as fine and true since they are committed to both conjuncts in each case being true. Consider for instance, if we substitute the generic conjunct with Leslie's weak truth-conditions or the capacity reading, respectively:

(33) a. Mosquitoes carry the West Nile virus, but typically they don't.
b. Some mosquitoes carry the West Nile virus and all of them are disposed to carry it, but typically they don't carry the West Nile virus.
c. Mosquitoes can carry the West Nile virus, but typically they don't.

The glosses of Leslie's weak truth-conditions and the capacity reading provided in (33b)-(33c) sound fine and are perfectly compatible with the claim that mosquitoes typically carry the West Nile. Thus, none of the competing proposals has an explanation of why, if their proposal is the preferred interpretation of (4), the conjunction in (33a) sounds perplexing in some way. This is problematic. The infelicity of the conjunctions in (30)-(32) indicates that (4)-(6) are in fact false, not some true weak reading.

An error theory can easily explain the contradictoriness or infelicity of (30)-(32). If we accept an error theory, then we can say that the first conjunct of (30a) is close in meaning to that of the negation of the second conjunct, and so their conjunction is false — i.e., the first conjunct is incompatible with the second. Our initial reaction that (4)-(6) are true is mistaken. (30)-(32) make the falsity of (4)-(6) plain and evident to us.

In making the above argument, I have used the vague expression "close in meaning" and have not elaborated on what I take the strong interpretation of the generic conjunct to be. I won't commit myself to a precisification of "close in meaning" or to any particular content for the generic conjunct, rather I will simply note that the argument above can be made precise in various ways. First, the argument is compatible with treating the generic

conjunct, in (30) for example, as close in meaning to *It's not the case that typically mosquitoes don't carry the West Nile virus*, and also as close in meaning to *Typically mosquitoes carry the West Nile virus*. Either treatment predicts that (30) sounds contradictory (though, note that the former might be consistent with *It's not the case that typically mosquitoes carry the West Nile virus*. Second, a more radical option to "close in meaning" is to say that the generic conjunct entails the other conjunct. In this case the generic conjunct is stronger and so not merely close in meaning. This stronger treatment also explains the contradictory feeling.

I will now address two worries about the evidence.

Evidence 1: Two Worries

Worry 1: The first worry is that the conjunction test over-predicts since one can purportedly get similarly infelicitous conjunctions when a generic is conjoined with an existentially quantified sentence — consider for instance:

(34) a. Ravens are black, but some aren't.
b. Some ravens are white, but ravens are black.
c. Ravens are black, but sometimes they're not.

Some of my informants claim that (34a)-(34c) sound perplexing, maybe even in a similar way to how (30)-(32) do. This would be problematic since, then, these conjunctions would indicate that (4)-(6) are close in meaning to a weak interpretation. I do not think that they sound equally perplexing or even perplexing in a similar way (i.e., a contradictory way). Here is a similar case which sounds fine:

(35) Ravens are black, but there are a few exceptions.

I contend, that if (35) is fine, then (34a)-(34c) should be fine as well.

If one still insists on hearing some infelicity in (34a)-(34c), then I claim that this is due to a tension in the relevance of the *some*-claim or the domains being talked about. I can hear a similar tension (to varying degrees) in explicitly quantified sentences, as below:

(36) a. Normally, ravens are black, but some aren't.
b. Some ravens aren't black, but typically ravens are black.
c. Generally, ravens are black, but sometimes they're not.

If the same tension can be generated with explicitly quantified sentences as in (36), then the conjunctions do not entail that the explicitly quantified sentences in question are close in meaning to their existentially-quantified counterparts.

Further, note that, as above, there are various ways of precisifying how the generic conjunct is close in meaning to the other conjunct. So even on the assumption that these cases do sound perplexing in some way (which I do not think that they do), there is a precisification which is compatible with treating the generic conjunct as having a strong interpretation: Take, for example, the generic conjunct of (34a), *Ravens are black*. It can

be given a strong interpretation if we treat it as stronger than (as entailing) *It's not the case that some ravens aren't black* — i.e., *All ravens are black.* Further, paired with domain restriction, this option is better than concluding that *Ravens are black* is as weak as an existential, as would be the case on the wide-scope reading of the second conjunct in (34a).

Worry 2: The second worry is that the conjunction test over-predicts with respect to what might be thought of as legitimately weak interpretations of generics. There are other generics, that seem very different in kind from (4)-(8), that could be thought of as having legitimately weak interpretations. If the test applies equally well to these, then the error-theorist needs to claim that we are making a mistake in these cases as well. But this seems wrong because our hypothesis is that they are legitimately weak interpretations. Consider the following:

(37) Birds lay eggs.

(38) Bees reproduce.

(39) Dobermans have floppy ears.

(40) Marine turtles are long-lived.

(41) Frenchmen eat horse meat.

(42) Scots wear kilts.

Each of these generics is intuitively true despite only a minority of the kind in question satisfying the predicated property: Only the adult female birds lay eggs; there is normally only one adult, mated queen bee that reproduces for each bee hive; Dobermans are born with floppy ears but breeders chop them off so that most have short pointy ears; although female marine turtles lay about 5000 eggs in their lifetime, it is unlikely that more than four of them will survive to adulthood; few Frenchmen have horse meat as a staple in their diet, nor would they eat it on special occasions, and finally, generally Scottish men nowadays only wear kilts for formal and some sporting occasions.

The error-theorist does not, or at least may not, want to be committed to the claim that generics like (37)-(42) are false, but it looks like the arguments she used in favour of an error theory work in the case of (37)-(42) as well. Consider:

(43) a. Birds lay eggs, but typically they don't.
b. Typically birds don't lay eggs, but birds lay eggs.

(44) a. Bees reproduce, but generally they don't.
b. Generally bees don't reproduce, but bees reproduce.

(45) a. Dobermans have floppy ears, but typically they don't.
b. Typically dobermans don't have floppy ears, but dobermans have floppy ears.

(46) a. Marine turtles are long-lived, but typically they aren't.
b. Typically marine turtles aren't long-lived, but marine turtles are long-lived.

(47) a. Frenchmen eat horse meat, but typically they don't.
b. Typically Frenchmen don't eat horse meat, but Frenchmen eat horse meat.

(48) a. Scots wear kilts, but typically they don't.
b. Typically Scots don't wear kilts, but Scots wear kilts.

The conjunctions in (43)-(48) sound contradictory or at least infelicitous in some way. But it is far from clear that the error-theorist wants to claim that (37)-(42) are false.

The error-theorist could of course go all out and claim that (37)-(42) are false. She could, for instance, appeal to the idea that (41) and (42) sound like clichés or silly stereotypes, and that such silly generalisations are false. However, the error-theorist need not go this route.

If the error-theorist does not want to claim that (at least) some of these generics, (37)-(40) for example, are false, then she needs to say what makes the interpretation of (4)-(8) different from that of (37)-(40)? The easy answer is that our best theories of domain restriction apply straightforwardly to make (37)-(40) true (on a strong interpretation), while they fail to make generics like (4)-(8) true (on a strong interpretation).

If domain restriction applies straightforwardly to (37)-(40), then what appears to be a weak interpretation is in fact a strong interpretation suitably restricted. Moreover, provided domain restriction works for (37)-(39), the error-theorist has an explanation of the contradictory conjunctions in (43)-(46).

For example, (43a) and (43b) sound contradictory, because they *are* contradictory. According to a reasonably standard theory of domain restriction for quantifiers — in particular, adverbs of quantification — the domain of quantification is restricted by the maximal set defined by the disjunction of the presuppositions of the predicate.[16] On such an account, the first conjunct of (43a) says something akin to (49a) or (49b) and the second conjunct of (43a) says something akin to (50):

(49) a. All normal birds that extrude offspring, lay eggs.
b. Generally, birds that reproduce in some way, lay eggs.

(50) Typically, birds that reproduce in some way, don't lay eggs.

On this construal, the first conjunct of (43a) is true, while the second conjunct is false, which explains why their conjunction sounds contradictory — it sounds bad because it is false. Similar observations hold for (44)-(46). Thus, the error-theorist has a plausible explanation of the contradictory conjunction data for generics that exhibit seemingly legitimate weak truth-conditions.

2.2 Evidence 2: Disagreement and Retraction

Additional evidence that we should not take our initial reaction to generics like (4)-(8) at face value comes from examples involving disagreement and retraction. Consider:

[16] See, e.g., A. Cohen (2008), Beaver and Clark (2008), Krifka (1995) and Rooth (1985, 1995).

(51) A: Let's stay inside. Mosquitoes are out there and they carry the West Nile.
B: That's ridiculous. Almost none of them do.

(52) A: Sharks attack bathers.
B: No, they almost never do.

(53) A: Rottweilers maul children.
B: Don't be silly. There have only been a few isolated incidences.

The dialogues in (51)-(53) appear to be genuine disagreements. Further, if they are indeed genuine disagreements, then we have evidence that generics like (4)-(8) are false and hence, for an error theory. To illustrate, in (51) for example, speaker B believes that the fact that almost no mosquitoes carry the West Nile is good reason for her to disagree with A's assertion of (4). Since B's assertion — that almost no mosquitoes carry the West Nile virus — is true, and A and B are disagreeing, this is evidence that (4) is false.

To bolster the argument, consider how one would expect speaker A to retract her statement. In examples like (51)-(53), speaker A would naturally retreat to something like: *Well, at least some of them do.* A natural explanation of this retraction is that A's initial statement — i.e., the generic claim — is false. These kinds of dialogues I suggest should be taken as evidence that (4)-(8) are false, and hence, for an error-theory.

Competing accounts of (4)-(8) cannot classify these dialogues as genuine disagreements and retractions (at least not without a substantive amount of work). Consider the dialogues with Leslie's weak truth-conditions and the capacity reading, respectively, substituted for the generic:

(54) A: Let's stay inside. Mosquitoes are out there. Some of them carry the West Nile and all of them are disposed to carry the West Nile.
B: That's ridiculous. Almost none of them do.
A: Well, at least some of them do.

(55) A: Let's stay inside. Mosquitoes are out there and they can carry the West Nile.
B: That's ridiculous. Almost none of them do.
A: Well, at least some of them do.

A's assertion that some mosquitoes carry the West Nile is obviously not incompatible with B's assertion that almost none of them do. Moreover, A's claim that all mosquitoes are disposed to carry the West Nile is not incompatible with B's claim that almost none of them in fact carry the disease. Both A's and B's claims are perfectly compatible and in fact true. Thus, A and B are not disagreeing in (54). A's second claim in (54) does not count as a retraction either: A's assertion that at least some mosquitoes carry the West Nile is compatible with both her original claims — that at least some mosquitoes carry the West Nile is only a slightly stronger version of her original statement that some mosquitoes carry the West Nile. Similar observations hold in the case of (55): That mosquitoes can carry the West Nile is perfectly compatible with the claims that almost none of them do and that at least some of them do. Thus, the competing proposals for (4)-(8) cannot, at least

straightforwardly, account for the disagreement and retraction evidence.

Evidence 2: Two Worries

Worry 1: One might have similar concerns to those raised in the previous section — that my diagnosis of the data over-predicts in various ways. It seems, for instance, that the disagreement data over-predicts since it is possible for disagreements to take place where the relevant statement is as weak as an existentially quantified statement. Consider, for instance:

(56) A: Sharks attack bathers.
B: No, some are perfectly gentle.

This seems problematic for anyone who wants to take disagreements like (51)-(53) as evidence for a strong interpretation of generics. However, my reaction to cases such as these is the same as it was for the contradictory conjunction cases above. Consider for example:

(57) A: Sharks normally attack bathers.
B: No, some are perfectly gentle.

(58) A: Sharks always attack bathers.
B: No, some are perfectly gentle.

If the same sense of disagreement can be generated with explicitly quantified sentences as in (57)-(58), then this obviously does not entail that the explicitly quantified sentences in question are close in meaning to their existentially-quantified counterparts. Thus, I contend, disagreement dialogues like the one in (56) do not entail that generics are close in meaning to existentially-quantified sentences or some suitably weak generalisation.

Worry 2: What of generics like (37)-(40), however? Again, the disagreement test can be thought to over-predict with respect to generics like (37)-(40). Consider, for instance:

(59) A: Let's get a Doberman. Dobermans have floppy ears.
B: That's ridiculous. Almost none of them do.

My reply, again, is the same as for the conjunction test above. The disagreement data in (59) can be dealt with using a standard treatment of domain restriction. In (59), the sense of disagreement results from A making a mistaken claim which B is correcting by pointing out that what A said is false and why it is false. According to our best theories of domain restriction, topic constrains the selection of the domain of quantification. Thus, in the case of (59) for instance, if the topic of the dialogue is pet Dobermans, then the domain picked out by A's statement is not the floppy eared Dobermans, but the pointy eared ones, so what A says is false. B points this out. The error-theorist has an explanation of disagreements involving generics which exhibit seemingly weak truth-conditions.

The difference between cases like (4)-(8) and those like (37)-(40) is that standard theories of domain restriction explain their seeming weakness.

As a final argument in favour of an error theory, I will propose a variant of Kripke's famous test, and argue that (4)-(8) do not count as counter-examples to the standard semantic analysis of generics, which treats them as having a strong interpretation.

2.3 Kripke's Test and Weak Interpretations of Generics

In 'Speaker's reference and semantic reference', Kripke recommends a test to determine whether alleged counterexamples to a semantic proposal should count as genuine counterexamples. According to Kripke's test:

> If someone alleges that a certain linguistic phenomenon in English is a counterexample to a given analysis, consider a hypothetical language which (as much as possible) is like English, except that the analysis is stipulated to be correct. Imagine such a hypothetical language introduced into a community and spoken by it. If the phenomenon in question would still arise in a community which spoke such a hypothetical language (which may not be English), then the fact that it arises in English cannot disprove the hypothesis that the analysis is correct for English. (1977, p.16)

In this context, something like Kripke's test can be a helpful heuristic for thinking about the examples at issue.[17] Kripke's original test was perhaps not intended for these purposes, but a variant can be used. The thought is that a semantic theory that treats our (initial) positive judgements about (some) weak generics as mistaken is supported by Kripke's test: In the language where we stipulate that such generics are false and speakers' psychological propensities remain as now, the initial (now clearly mistaken) intuition is still likely to be present.

The idea is that this variant of Kripke's test can be used in favour of the conclusion that cases like (4)-(8) do not count as counterexamples against the hypothesis that they receive a strong interpretation, i.e., the same interpretations as paradigmatic generics like (1)-(3). In other words,(4)-(8) needn't be given a weak interpretation whereby they are only true in virtue of a minority of the instances in question satisfying the predicated property.

Applying this variant of Kripke's test, suppose that L is a language in which generics like (4)-(8) are given an analysis consistent with the strong interpretation — e.g., for (4), something along the lines of *all normal mosquitoes carry the West Nile Virus* or *generally, mosquitoes carry the West Nile Virus.* Further stipulate that L is a language just like English in

[17] This isn't the place for an evaluation of Kripke's test. It is generally taken to be at least a helpful heuristic for determining whether a given linguistic phenomenon is semantic content. That said, it is worth noting that there is disagreement about just what the test is a diagnostic of. Kripke takes it to provide support for semantic analyses (facing a set of counterexamples), but it takes argument to show that a given analysis is a correct account of the semantics, rather than some other level of content. I will not try to settle that issue here.

every other respect and that the psychology of speakers of L is as close as possible to the psychology of speakers of English. Now we ask the question: Would speakers of L make the mistake of thinking that (4)-(8) are true? Would a minority of the kind in question having some salient striking property lead speakers to make a mistaken generalisation?

If your inclination is to answer yes to these questions, then this is no argument against English being just like L — i.e., English having a semantics for cases like (4)-(8) on which they are given an analysis consistent with the strong interpretation. This is further support for an error theory.

I have now presented three reasons to prefer an error theory to deal with the (4)-(8) cases. In the next section, I use the (4)-(8) examples as a case study to provide some more general conclusions about when error-theories are needed in semantics.

3 Content, Interpretation and Cognitive Bias

When it appears as though our judgments about the truth-conditions of sentences are somehow sensitive to cognitive bias, such as our tendency to overgeneralise when the predicated property is somehow striking or appalling, there are several options for how to incorporate the impact such a bias on the truth-conditions: As part of the content of some expression in the utterance (as in for example, evaluative or response-dependent terms), as part of the pragmatically derived content, as part of the metasemantics of context-sensitive terms (as in for example, the selection of a salient quantifier domain or salient referent of *he*), or finally, as no part of the communicated content at all. The last option is what I have called an "error theory".

It is worth pointing out that there are interesting non-generic analogues to cases like (4)-(8), which bring up the same issues. In principle, a large chunk of empirical psychology is devoted to errors in our judgments of truth and many of these will be relevant to the intuitions we appeal to as evidence in our semantic theories. As an example: It is well-known that probability judgements are subject to certain cognitive biases — for example, the *representativeness heuristic*, which is the tendency to judge an event or hypothesis as more probable by considering how much the event or hypothesis resembles available data as opposed to using the probability calculus as a basis for the judgement — it would be interesting to see the impact of such features on our accounts of the truth-conditions of sentences containing probability operators. Suppose we have semantic data analogous to the famous Linda-the-bank-the-teller cases of Tversky and Kahneman (1983), and hence, that speakers intuitions about the truth-conditions of sentences containing probability operators inherit in some way the pitfalls of the conjunction fallacy. Should we then automatically trust our intuitions and assume a semantics of probability operators that invalidates such inference patterns as the distribution of probabilities over conjunction? As far as I'm aware, in this case no one has suggested that we adjust the semantics in a way that makes the biases part of the truth-conditions of probability operators. If our probability talk and thought does not at least in some significant way give voice to the

probability calculus, then this would be at least prima facie problematic for various reasons. Given that the probability calculus is the most objective and logical means we have to track and voice prevalence and likelihoods, one would hope that is what our probability talk is about in some significant sense.

Several authors have proposed error theories as viable options along side semantic and pragmatic proposals of problematic cases.[18] This should be done more often. An interesting and well developed case-study is in the "stakes-sensitivity" of knowledge ascriptions. Theorists have been ardently debating whether or not *knows* is *stakes-sensitive*. According to many influential epistemologists,[19] when presented with scenarios of differing stakes, our intuitions about the the truth-value of a given knowledge ascription vary. An influential position, *epistemic contextualism*, takes these intuitions to indicate that the truth-conditions of *knows* include reference to the stakes. By contrast, Williamson (2005) and Nagel (2008) argue that knowledge ascriptions are invariant despite intuitions to the contrary, and what accounts for these intuitions is cognitive bias — for Williamson, the error is due to a failure to recognise the distinction between knowing and knowing that one knows, while for Nagel the error is due to a cognitive factor called *need for closure*.[20] What Williamson and Nagel provide are error theories, along similar lines to what I am proposing for the (4)-(8) cases.

The point of these case studies is to show that in general, it is a mistake to assume that cognitive biases or effects on our truth-value intuitions for a class of sentences, S, should be predicted by a semantics for S. In general, when we are in a domain where we know that psychological biases are prevalent, the first (but not the last) option to consider is that these biases generate mistakes. A plausible methodological principle is the following: 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.

References

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

Asher, N., & Pelletier, F. (2012). More truths about generic truth. In A. Mari, C. Beyssade, & F. D. Prete (Eds.), *Genericity*. Oxford: Oxford University Press.

Beaver, D., & Clark, B. (2008). *Sense and sensitivity: How focus determines meaning*. Oxford: Wiley-Blackwell.

Bhatt, R. (1999). Ability modals and their actuality entailments. In K. Shahin, S. Blake, & E.-S. Kim (Eds.), *WCCFL 17 proceedings*. Stanford: CSLI Publications.

[18] For example, Braun and Saul (2002) provide a well developed error theory for conflicting intuitions about substitutions into simple sentences.

[19] See, e.g., DeRose (1992), S. Cohen (1999), Hawthorne (2004) and Stanley (2005).

[20] See, e.g., Kruglanski (1989) and Kruglanski and Webster (1996).

Braun, D., & Saul, J. (2002). Simple sentences, substitutions, and mistaken evaluations. *Philosophical Studies*, *111*, 1-41.

Carlson, G. (1977). *Reference to kinds in english.* Amherst: Doctoral Dissertation, Department of Linguistics, University of Massachusetts Amherst.

Carlson, G. (1989). The semantic composition of english generic sentences. In G. Chierchia, B. Partee, & R. Turner (Eds.), *Properties, types and meanings (vol. 2): Semantic issues.* Dordrecht: Kluwer.

Cimpian, A., Brandone, A. C., & Gelman, S. A. (2010). Generic statements require little evidence for acceptance but have powerful implications. *Cognitive Science*, *34*, 1452-82.

Cohen, A. (1996). *Think generic!: The meaning and use of generic sentences.* Pittsburgh: Doctoral Dissertation, Department of Linguistics, Carnegie Mellon University.

Cohen, A. (2001). Relative readings of many, often and generics. *Natural Language Semantics*, *9*, 41-67.

Cohen, A. (2008). No alternative to alternatives. *Journal of Semantics*, *26(1)*, 1-48.

Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, *13*, 57-89.

DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, *52*, 913-29.

Hawthorne, J. (2004). *Knowledge and lotteries.* Oxford: Oxford University Press.

Kratzer, A. (2012). *Modals and conditionals: New and revised perspectives.* Oxford: Oxford University Press.

Krifka, M. (1995). Focus and the interpretation of generic sentences. *The Generic Book*, 238-64.

Kripke, S. (1977). Speaker's reference and semantic reference. In P. French, T. Uehling, & H. Wettstein (Eds.), *Contemporary perspectives in the philosophy of language.* University of Minnesota Press.

Kruglanski, A. (1989). *Lay epistemics and human knowledge.* New York: Plenum Press.

Kruglanski, A., & Webster, D. (1996). Motivated closing of the mind: "seizing" and "freezing". *Psychological Review*, *103*, 263-83-.

Leslie, S. (2007a). Generics and the structure of the mind. *Philosophical Perspectives*, *21(1)*, 375-403.

Leslie, S. (2007b). *Generics, cognition and comprehension.* Princeton: Doctoral Dissertation, Department of Philosophy, Princeton University.

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

Leslie, S. (2013). Generics oversimplified. *Nous*.

Liebesman, D. (2011). Simple generics. *Nous*, *45(3)*, 409-42.

Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, *86*, 279-94.

Nickel, B. (ms.a). *Between logic and the world: An integrated theory of generics.* Cambridge, Massachussetts: Draft Manuscript, Department of Philosophy, Harvard University.

Nickel, B. (ms.b). Mosquitoes and sharks, habit and capacities. *Draft, Department of*

Philosophy, Harvard University.

Rooth, M. (1985). *Association with focus.* Amherst: Doctoral Dissertation, Department of Linguistics, University of Massachusetts Amherst.

Rooth, M. (1995). Indefinites, adverbs of quantification and focus semantics. In G. Carlson & J. Pelletier (Eds.), *The generic book.* Chicago: University of Chicago Press.

Rothbart, M., Fulero, S., Jenson, C., Howard, J., & Birrell, P. (1978). From individual to group impressions: Availability heuristics in stereotype formation. *Journal of Experimental Social Psychology, 14,* 237-55.

Schaffer, J., & Szabo, Z. (2014). Epistemic comparativism: A contextualist semantics for knowledge ascriptions. *Philosophical Studies, 168(2),* 491-543.

Stanley, J. (2005). *Knowledge and practical interests.* Oxford: Oxford University Press.

Sterken, R. (forthcoming). Leslie on generics. *Philosophical Studies.*

Sterken, R. (ms.). Generics, covert structure and logical form. *Draft, University of Oslo.*

Tversky, A., & Kahneman, D. (1983). Extension versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90 (4),* 293-315.

Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *Philosophical Quarterly, 55,* 213-35.