



Taking truth seriously: the case of generics

Martin Gustafsson¹ 

Received: 26 September 2022 / Accepted: 29 May 2023 / Published online: 21 June 2023
© The Author(s) 2023

Abstract

By discussing a large number of different examples, this paper argues that the class of so-called generic statements is much more heterogeneous than is usually recognized in the contemporary debate. It is claimed that the theoretical tendency towards overgeneralization or homogenization makes it impossible to adequately understand how generic statements function in language and to handle the dangers involved in generics that express and promote social stereotypes and prejudices. It is also argued that such overgeneralization involves what J. L. Austin described as the “fetishizing” of the true/false distinction. Several influential theories are criticized, including Sterken’s error theory, Leslie’s conception of weak generics, and the idea that generics involve loose talk.

Keywords J. L. Austin · Generics · Leslie · Stereotype · Sterken · Truth

Difference is resource, not failure
Helen E. Longino (1994, p. 477).

In lecture XII of *How to Do Things with Words*, J. L. Austin famously admits to an “inclination to play Old Harry” with “the true/false fetish” (Austin, 1975, p. 151). It is a much discussed issue what, exactly, Austin has in mind here. It might be tempting to think that he somehow wants to devalue or discredit truth, or at least show that philosophers have overrated its significance.

However, making a fetish out of something is not the same as overrating its significance. Rather, fetishizing might involve misconstruing and thereby losing sight of the thing’s real importance. Austin thinks philosophers have disengaged the true/false distinction from the concrete work it performs in real-life circumstances of language use, but he nowhere suggests that this work would be unimportant. Arguably, his point is the very opposite: the significance of the true/false distinction consist precisely in this work, and thus the philosophical tendency to discuss truth in abstraction from

✉ Martin Gustafsson
martgust@abo.fi

¹ Philosophy/FHPT, Åbo Akademi University, Arken, Fabriksgatan 2, 20500 Turku, Finland

real-life use in effect deprives the distinction of its importance. Philosophers *think* that they take truth seriously, but in fact they disconnect truth from what makes it worthy of our esteem in the first place.

According to Austin, this disconnection is manifested in the philosophical tendency to think that the question ‘True or false?’ is applicable to any constative utterance, merely *qua* constative. However, he notes, “[i]n real life, as opposed to the simple situations envisaged in logical theory, one cannot always answer in a simple manner whether [a constative utterance] is true or false” (Austin, 1975, p. 143). The first example he uses to illustrate his point has become a classic, and his discussion is worth quoting at length:

Suppose that we confront ‘France is hexagonal’ with the facts, in this case, I suppose, with France, is it true or false? Well, if you like, up to a point; of course I can see what you mean by saying that it is true for certain intents and purposes. It is good enough for a top-ranking general, perhaps, but not for a geographer. ‘Naturally it is pretty rough’, we should say, ‘and pretty good as a pretty rough statement’. But then someone says: ‘But is it true or false? I don’t mind whether it is rough or not; of course it’s rough, but it has to be true or false – it is a statement, isn’t it?’ How can one answer this question, whether it is true or false that France is hexagonal? It is just rough, and that is the right and final answer to the question of the relation between ‘France is hexagonal’ and France. It is a rough description; it is not a true or false one. (Austin, 1975, p. 143)

Notice that Austin is not just making the sort of point familiar from contemporary contextualist debates, that the truth-conditions of ‘France is hexagonal’ depends on the particular circumstances of use. He is *not* saying that the sentence is straightforwardly true in the mouth of a top-ranking general and straightforwardly false in the mouth of a geographer. What he is saying is that the sentence is “good enough” for the general, and in this sense “true up to a point” or “true for certain intents and purposes”—but in saying so, he is questioning the idea that the true/false distinction is straightforwardly applicable *even* when it has been determined that the sentence is uttered in these military-strategic circumstances. The right and final characterization, he says, is that the description is rough—too rough for the geographer, but good enough for the general. According to Austin, the use of the true/false distinction is not just more context-sensitive than philosophers tend to assume, but in a sense more *specialized*: its significance comes out in the fact that in many cases, *other* terms of assessment are more appropriate even if what we have before us is a declarative sentence uttered in a specific context.

Austin’s second example is less famous, but equally interesting:

Consider the constative, ‘Lord Raglan won the battle of Alma’, remembering that Alma was a soldier’s battle if ever there was one and that Lord Raglan’s orders were never transmitted to some of his subordinates. Did Lord Raglan then win the battle of Alma or did he not? Of course in some contexts, perhaps in a school book, it would be perfectly justifiable to say so – it is something of an exaggeration, maybe, and there would be no question of giving Raglan a medal for it. As ‘France is hexagonal’ is rough, so ‘Lord Raglan won the battle of

Alma' is exaggerated and suitable to some contexts and not to others; it would be pointless to insist on its truth or falsity. (Austin, 1975, pp. 143–144)

Again, Austin's main claim is not that the truth-conditions of 'Lord Raglan won the battle of Alma' vary with the context of utterance. Rather, it is that insisting on the statement's truth or falsity would be misguided, even in certain specific contexts of use. Such insistence would be pointless, since more nuanced descriptions that do not use the words 'true' and 'false' give a better and more precise account of the relation between the statement and reality: it might be perfectly justifiable in a schoolbook to say that Lord Raglan won the battle, but something of an exaggeration, and so on and so forth. Again, what Austin points out is that the significance of the true/false distinction is tied to its having a more specialized use than philosophers tend to assume, and, hence, that insisting on its general applicability means losing sight of how the distinction actually matters to us.

My aim in this paper is not to engage in Austin exegesis. Instead, I investigate what Austin's point amounts to by applying it to a hot topic in contemporary philosophy: the nature and danger of so-called generic statements (or 'generics' for short). By discussing a large number of different examples, I argue that much of the contemporary debate about generics suffers from fetishizing the true/false distinction. Such fetishizing makes it impossible to adequately understand how generics function in language and to handle the dangers involved in generics that express and promote social stereotypes and prejudices.

1 Introductory ornithological exercises

Let's warm up with the following example:

(1) Birds fly.

This sentence is often used in discussions of generics, and is usually said to be true (cf., for example, Cohen, 2012; Thakral, 2018). This may seem innocuous enough but is in fact contestable. Consult any book on bird zoology: nowhere will flying be listed among the characteristic features of the group of vertebrates that constitute the class *Aves*, with its almost 10,000 species. Rather, you will learn that birds have feathers, toothless beaked jaws, a four-chambered heart, a strong but lightweight skeleton, and that they lay hard-shelled eggs. There are more than 60 extant species of bird that do not fly, and many more extinct ones. All ratites are flightless (they include ostriches, emus, rheas, cassowaries, kiwis, and others), as are all members of the penguin family. Some duck and rail species do not fly, and then there is the flightless owl parrot or kakapo.

But isn't that precisely the thing about generics—that they tolerate exceptions and can therefore be true even in the face of apparent counterexamples? Indeed, 60 out of nearly 10,000 species seem pretty marginal, considering that there are true generics which allow for a *much* higher proportion of exceptions (more on such so-called 'weak' generics in Sect. 2). So what's the problem? Isn't 'Birds fly' in fact a very good and representative example of a true generic statement?

No. Again, no book on bird zoology would list flying among the characteristic features of members of the *Aves* class. In situations in which truth is genuinely at stake, the statement ‘Birds fly’ requires qualification. The statement is rough and oversimplified. Why?

It has to do with the nature of the exceptions. There is nothing the matter with an ostrich that does not fly. That’s what an ostrich is like, *qua* ostrich. And being an ostrich is a perfectly fine way of being a bird—as is being an emu, penguin, kiwi, or kakapo. ‘Birds fly’ is a sloppy generalization to which people who are knowledgeable about birds do not assent without qualification.

By contrast, consider what happens if we descend the taxonomic ladder a bit. The following sentence is taken from the Wikipedia entry on flamingos:

(2) Flamingos are capable flyers.

This is a good example of a true generic. The exceptions are there, of course: Flamingos that cannot fly due to birth defects or wounds, and baby flamingos that have not yet learnt to fly. But these exceptions do not undermine the truth of (2), for they are quite different from the ones considered above in relation to ‘Birds fly’. Unlike ostriches and members of other non-flying species, individual flightless flamingos are either incapacitated or immature. Indeed, their being either incapacitated or immature reflects the truth of (2): *since* flamingos fly, an individual flamingo’s inability to fly constitutes an defect or immaturity.¹

It may be objected that what I have said about ‘Birds fly’ is correct only with regard to a narrow context with very strict standards—viz., the scientific context of bird zoology. Arguably, there are other contexts in which the standards are more relaxed, and where the statement ‘Birds fly’ is true. For example, it would be obviously misplaced to object against Oscar Hammerstein II’s immortal lines, ‘Fish gotta swim, birds gotta fly, I gotta love one man ‘til I die’, by pointing out that penguins do not fly. And here is Matt Teichman:

[I]magine a parent who is trying to teach their two-year-old child what a bird is. Any good explanation of what a bird is should include the information that birds fly. A parent who told their child that birds do not fly would be actively shirking their parental responsibilities. (Teichman, 2015, p. 9)

This objection, however, is built on a misunderstanding. To begin with, my point is not that it is *false* to say that birds fly, or that its negation, ‘Birds do not fly’ is true. Nor am I arguing that ‘Birds fly’ is false in some contexts and true in others. Rather, my claim is the Austinian one, that ‘Birds fly’ is a rough statement that may be good enough (though not straightforwardly true) in some situations but too crude (though not straightforwardly false) in others. It *is* good enough in Hammerstein’s lyrics, and good

¹ What about the true generic ‘Flamingos lay eggs’? After all, male flamingos lay no eggs, and yet they are perfectly fine *qua* flamingos. (Similarly, ‘Birds lay eggs’ is true even if male birds are perfectly fine *qua* birds). My response is that this is an exceptional case: since an organism is by definition female if it produces the ovum, ‘Flamingos lay eggs’ and ‘Birds lay eggs’ patently purport to speak only of female birds/flamingos. Contrast this with, say, ‘Pheasants are highly decorated with bright colors and adornments such as wattles’, which does require qualification since it happens to be true only of male pheasants. Thanks to an anonymous reviewer for asking me to consider this objection.

enough for the parent in Teichman's example. And of course, it is much better than its negation—Teichman is right that telling one's child that birds do not fly amounts to shirking one's parental responsibilities. However, this does not entail that 'Birds fly' is *true*, even in the mouth of the imagined parent. The situation Teichman describes is one in which truth is not insisted upon, and for good reasons: a two-year-old is not yet ready to engage in proper ornithological truth-seeking. The parent's use of rough statements such as 'Birds fly' is better seen as a preparatory step, leading the child on toward a more mature state in which she will be able to appreciate what it takes to distinguish between ornithological truths and falsehoods. As such, the statement is perfectly legitimate, and, indeed, incredibly important. However, it is not true (and not false).

Another objection proceeds from the observation that sentences such as the following (taken from the online version of the *Encyclopedia Britannica*) are undeniably true:

(3) Birds fly by flapping their wings.

Arguably, (3) entails 'Birds fly', so mustn't we acknowledge the truth of (1)?² On further reflection, however, the alleged entailment seems spurious. For what does (3) do? It provides an explanation of the phenomenon of bird flight. Consequently, its scope of reference is patently restricted to species of bird that *do* fly. If someone objects, 'What about penguins?', the utterer of (3) can sensibly answer, 'Don't be stupid, of course the explanation only deals with birds that fly and not with cases in which there is nothing for it to explain.' A substantive objection against (3) would instead invoke species of bird that do fly but do it in some other way than by flapping their wings (in fact there are no such species of bird; some large birds such as eagles sometimes fly by gliding, but they also fly by flapping their wings).³

What can be learnt from our warming-up exercise? An immediate lesson is this: the generics 'Birds fly' and 'Flamingos fly' differ due to the fact that 'Birds' refers to a

² Thanks to an anonymous referee for raising this objection, and for the *Encyclopedia Britannica* reference.

³ Note three things. First, I do not deny that 'Birds fly by flapping their wings' entails 'Some birds fly'. Second, I can even afford to admit that there is a possible reading of 'Birds fly by flapping their wings' according to which it *does* entail the self-standing generic 'Birds fly'. I take such a reading to be unnatural in most contexts, but the crucial point is that *if* 'Birds fly by flapping their wings' is thus construed, it is no longer straightforwardly true. Rather, it is just as rough and unqualified as 'Birds fly' (and similarly vulnerable to the objection 'What about penguins?'). So, my diagnosis is this: When we take 'Birds fly by flapping their wings' to be patently true, we tacitly rely on a reading according to which it does *not* entail the self-standing generic 'Birds fly' (but rather something like the standardly quantified 'Some birds fly'). However, when we then go on to argue that 'Birds fly by flapping their wings' does entail the self-standing generic 'Birds fly' – tacitly and mistakenly taking it for granted that just because the *word-string* 'birds fly' occur in 'Birds fly by flapping their wings', this entailment relation must subsist – we abandon precisely the sort of reading on which we relied when we called 'Birds fly by flapping their wings' true in the first place. Third, one may wonder: If 'Birds fly by flapping their wings' is restricted in scope to flying species of bird, then why cannot the isolated generic 'Birds fly' itself be read as restricted in this way and thus as in effect synonymous to 'Flying birds fly' (which is tautologically true)? My answer is that in the case of 'Birds fly' such a restriction would construe the sentence as not even purporting to be informative, which goes against any reasonable strategy of interpretation (including Grice's maxim of quantity). The same point applies in response to the worry that if 'Flamingos lay eggs' is restricted in scope to female flamingos, then why cannot 'Flamingos are female' be heard as similarly restricted in scope and thus be true? Unlike 'Female flamingos lay eggs' which is informative (since it excludes, for example, that female flamingos give birth to live young), 'Female flamingos are female' does not even purport to be informative. Hence the envisaged restriction of scope is excluded in the latter case.

whole class of species, whereas ‘Flamingos’ refers to a single species. This difference comes out in the different logics of the two statements. More precisely, ‘Birds fly’ differs from ‘Flamingos fly’ in that one can object against it by invoking flightless species of bird as counterexamples. Such an objection is a call for qualification: what it shows is that ‘Birds fly’ is too rough and sweeping to be simply true. ‘Birds fly’ would be true if there were no flightless species of bird – just as ‘Birds have feathers’ *is* true, since there are no featherless species of bird (even if many pictures of poor featherless bird individuals can be found on the internet). In that sense, ‘Birds fly’ is a sort of generalization over all species of bird, and a single flightless species suffices to undermine its truth. However, unlike a universal generalization in standard quantificational logic, its truth is then undermined, *not* in the sense that the generic gets *falsified* – for of course, it isn’t false that birds fly! Rather, what the counterexamples show is that the generic is *roughly but not quite true*. (And the more counterexamples there are, the rougher the generic is; if the counterexamples are very many, a generic of this sort may even be too rough to be roughly *true*—which is not to say that it is false (or even roughly false).)

There is nothing particularly strange or obscure about this usage. However, if one is held captive by the idea that declarative sentences, merely *qua* declarative, must be either true or false, one will be blind to its very possibility. It will then seem as if one must classify ‘Birds fly’ either as true or as false. And since it seems very strange to call it false, the most natural choice will be to classify it as true. But how *can* it be true, if there are flightless species of bird? At this point, two alternatives may seem tempting. Either, one may try to argue that ‘Birds fly’ is correctly analyzed in non-generic, quantificational terms such as ‘Some birds fly’, or perhaps ‘Most birds fly’—thus concluding that its generic form is a mere surface phenomenon. Or, one might try to think of it as similar to ‘Flamingos fly’, and thus conclude that the objection ‘What about penguins?’ (directed against ‘Birds fly’) is just as misguided as the objection ‘What about individual flamingos that have hurt their wings?’ (directed against ‘Flamingos fly’). In either case, the true/false fetish is at work: the presumption that declarative sentences must be either true or false makes one blind to significant differences in use, and thereby fosters the assimilation of what are in fact quite diverse forms of description. In what follows we will encounter and look in detail at many other examples of such assimilation, and I will explain why this homogenizing tendency should often be resisted.

2 Loose talk and weak generics

If I am right, it is no coincidence that ‘Birds fly’ is sometimes used in the literature to explain why it may seem attractive to treat generic statements as *loose talk*. Bernhard Nickel asks us to consider the exchange,

- A. Birds fly.
 - B. What about penguins?
 - A. #Birds fly.
- (Nickel, 2016, p. 28)

He comments:

A's simple repetition in this exchange is odd, and this oddity is easily accounted for if generics were simply instances of loose talk. On this strategy, generics have the same truth-conditions as the corresponding universal generalization, but the bare plural indicates that we are speaking somewhat loosely. A's initial assertion is thus false – but it might nonetheless be unobjectionable if the flightless birds can be ignored. B's question, and its concomitant introduction of penguins into the conversation, makes it so that they cannot be ignored any longer. Hence the falsity of A's utterance can no longer be overlooked, as shown by the unacceptability of repeating it. (Nickel, 2016, p. 29)

It should be clear that I do not subscribe to the idea that 'Birds fly' is an instance of loose talk in the sense explained by Nickel (this notion of loose talk stems from Lasersohn, 1999). My claim is not that 'Birds fly' has the truth-conditions of a universal generalization in standard quantificational logic and that it is therefore false though unobjectionable in situations where the exceptions can be ignored for pragmatic reasons. Rather, what I have argued is that 'Birds fly' is neither true nor false but rough, and that it is therefore unobjectionable in situations where such roughness can be tolerated or is even advantageous. In *that* quite different sense, however, I am willing to agree that 'Birds fly' is loose—in contrast to 'Flamingos fly', which is indeed true despite the existence of baby flamingos and flamingos with wounded wings.

Now, Nickel does not defend the idea that generics constitute loose talk in the Lasersohnian sense. In fact, he criticizes this idea by means of an example analogous to 'Flamingos fly', namely, 'Ravens are black'. Nickel points out that it is perfectly fine to say, 'Ravens are black, though some ravens are white'—albino ravens do not falsify the generic. Hence, he concludes, generics are not loose talk.

This line of reasoning is peculiar. Even if Nickel notices that 'Birds fly' and 'Ravens are black' are different in that the former is vulnerable to counterexamples in a way that the latter is not, he makes nothing of this difference. He does not even make any effort at assimilation along the lines I suggested at the end of Sect. 1 but ends up simply neglecting the difference that he has just identified. Why? Well, he takes himself to be addressing a question about generics *überhaupt*—Are generics instances of loose talk?—and he uses the raven example as evidence for a *general* rejection of the loose talk conception of generics. He nowhere considers the possibility that there really is a deep-going difference between 'Birds fly' and 'Ravens are black', even if his discussion of the two cases would seem to support precisely such a conclusion.

Again, we encounter the tendency to overgeneralize and thereby homogenize, common among theorists in this field. In what follows, I aim to swim against this tide, by emphasizing how important it often is to register and keep in mind differences between generic statements if we are to properly understand and assess their philosophical significance. Here is one case which illustrates this point in relation to the examples just discussed. Suppose we want to understand and assess Michael Thompson's idea that the realm of the living has a "form" or "logic" that is manifested in what he calls "natural-historical judgments" (Thompson, 2008). According to Thompson, such natural-historical judgments are generic in character. However, he is very careful *not* to count statements such as 'Birds fly' among the natural-historical judgments.

By contrast, ‘Flamingos fly’ is a natural-historical statement in Thompson’s sense, and the logic of life that he is concerned to identify is visible precisely in the characteristic ways in which such statements tolerate counter-examples—ways having to do specifically with biological development and capacity. I will not defend Thompson’s view; my point is just that it can be properly evaluated only if we notice that he is concerned with a fairly *specific* kind of generic statement whose character he describes with considerable care. No more general account of generic statements can be straightforwardly employed to defend or reject Thompson’s conception.

In fact, I will argue that generic statements constitute a *much* more heterogeneous class than is sometimes assumed in the debate. A wish for homogeneity, reinforced by the presumption that declarative sentences are either true or false and that other terms of assessment are of secondary importance, often makes theorists lose sight of philosophically significant differences between kinds of generic statement. Sometimes, such striving for theoretical homogeneity can perhaps be justified if the theorist is engaged in a technical project of semantic modelling, the limited claims of which are clearly stated. However, as I will argue in the final section of this paper, such an excuse is by no means always appropriate. Often, differences that are highly significant given the stated purpose of one’s investigation are ignored or papered over.

The debate over so-called ‘weak’ generics—generics that apparently tolerate *very* many exceptions—provides illuminating examples. Consider Rachel Sterken’s discussion of the following five statements:

- (4) Mosquitoes carry the West Nile virus.
- (5) Sharks attack bathers.
- (6) Rottweilers maul children.
- (7) Tigers eat people.
- (8) Australian snakes are poisonous.

According to Sterken, these examples are “intuitively true even though only a minority, even a very small minority, of the kind satisfies the predicated property” (Sterken, 2015, p. 75). This is a fairly typical assumption among participants in this debate. Most of these participants also take it for granted that a proper account of generics should tally with our intuitions; Sterken is unusual in proposing an error theory according to which (4)–(8) are all false, despite our alleged intuitions to the contrary. And yet, she shares two crucial premises with her opponents: that these sentences are all intuitively true, and that the question ‘Are they true or false?’ is well posed and has the same answer in all cases (according to Sterken, the answer is ‘false’; according to her opponents the answer is ‘true’). I shall question both these assumptions.

To begin with, it should be pointed out that the notion of being ‘intuitively true’ is quite obscure. Presumably, it is supposed to mean something like ‘rings true’, or ‘appears true at first sight’—but then, what is *that* supposed to mean? In claiming that all these sentences are intuitively true, are we making a prediction about how ordinary speakers would classify them if they were given a questionnaire in which they were asked to assign a determinate truth-value to each one of them without much prior reflection? Or are we making some sort of introspective report? But why think that such procedures would tell us anything important about the roles of these statements in real-life language use? After all, in cases in which there is

a genuine need for knowledge about mosquitoes, sharks, Rottweilers, tigers, or snakes, a speaker's merely relying on her 'intuitions' is hardly a manifestation of her linguistic competence. The thing to do would rather be to consider an encyclopedia or ask someone who really knows about how such creatures function. Mastering the concept of truth involves *not* being satisfied with what merely 'rings true' or 'seems true at first sight', so if we ask people to base their truth-value assessments on such 'intuitions' we're making them behave in contravention of a principle which is fundamental to the real-life significance of the true/false distinction.

Keeping this point in mind, let us now proceed by taking a closer look at each one of the above statements, without presupposing that they must all fit one theory or that we have and share determinate 'intuitions' about their truth-value. How about 'Mosquitoes carry the West Nile virus'—a favorite example in this debate? Well, Swedish mosquitoes don't, so even if mosquitoes are a real plague in many parts of the country, most Swedes have probably never heard about the West Nile virus. Virtually all Swedish mosquitoes belong to the *Aedes* genus, whereas the West Nile virus is mainly carried by mosquitoes in the genus *Culex*. Among the approximately 3600 species of mosquitoes that exist in the world, approximately 65 are known to carry the West Nile virus (Colpitts et al., 2012). It is unknown exactly which mosquito species can carry the virus if they are exposed to it.

Culex mosquitoes are common in the US, and in many parts the species *Culex tarsalis* is the main transmitter of the West Nile. In Mississippi, if someone told a Swedish visitor about local hazards and in this context said, 'Mosquitoes carry the West Nile virus', that would be a valuable piece of information. It would presumably be obvious to the hearer that the intended domain was restricted to mosquitoes in the area. However, in order to focus on the special character of weak generics, let us bypass questions about whether and how the extension of the term 'mosquitoes' may depend on the geographical location at which the utterance is made, and instead consider a refined version of (4):

(4*) In the US, *Culex tarsalis* mosquitoes carry the West Nile virus.

This is a fine example of a weak generic statement. It is true (or merely seems true, if Sterken is right), even if less than one percent of US *Culex tarsalis* mosquitoes actually carry the virus. How is this possible?

Sarah-Jane Leslie takes all the weak generics listed above to be what she calls "Type B generics".⁴ According to her, when it comes to Type B generics, it is relevant, indeed crucial, that the property at issue is *striking*, "often in virtue of being dangerous or appalling" (Leslie, 2008, p. 40). This seems to fit (3*): The West Nile virus is dangerous and thus striking, and if Leslie is right this is necessary for its truth (Leslie, 2008, p. 43).⁵ Leslie's account has been very influential, but it seems to me that it is vulnerable to straightforward counterexamples. Thus, imagine a virus, the XYZ virus, which only

⁴ She divides generics into two main types, Type A and Type B (she also calls Type B generics "striking property generics"). In a Type A generic, the property ascribed lies "along a characteristic dimension" for the kind talked about, or (if what is at issue an artifact or a social kind) constitutes the function or purpose of the kind. There are some weak generics of Type A, such as 'OrangeCrusher 2000s crushes oranges' (which is true even on the assumption that no OrangeCrusher 2000 has ever been used).

⁵ Leslie says she does not propose this as a semantically derived truth-condition, but as a "worldly truth-specification", which is a matter of "how the world must be for the sentence to be true" (Leslie 2008, p. 43).

causes very mild coldness symptoms among cottontails. It is completely harmless to humans. If anything is non-striking, it is this virus. And yet, its inconspicuousness would not seem to make the sort of difference that Leslie proposes. After all, it might well be discovered that the XYZ virus is carried by US *Culex tarsalis* mosquitoes, just as it has been discovered that they carry the West Nile.

If so, and even if the frequency of carriers proved to be no higher than in the West Nile case, the following statement would be just as true as (4*):

(4**) In the US, *Culex tarsalis* mosquitoes carry the XYZ virus.⁶

Now, Leslie is certainly right that we have a worrisome tendency to generalize from very few striking instances to sweeping and highly dubious claims which sometimes have a generic shape. I will come back to the significance of this tendency in the next section, where I discuss so-called social generics. However, the West Nile virus example is importantly different from such dubious generalizations. After all, it is not just a dubious generalization but perfectly true that *Culex tarsalis* mosquitoes in the US carry the West Nile virus.

Two further points should be made. First, as Sterken rightly points out against Leslie, the striking character of a property can at most explain why we spontaneously *take* a weak generic to be true. It has no bearing on whether it is *in fact* true. Second, in the case of the West Nile example, it is not so clear that the striking character of the property has such a significant influence even on our ‘intuitions’ about the truth-values of (4*) and (4**). Of course, it is plausible that the inconspicuous character of the XYZ virus means that we won’t notice it in everyday life—probably, its existence and transmission routes will be recognized only by cottontail experts. So, (4**) won’t be a truth that most of us know or care about. By contrast, since the West Nile virus is dangerous, we will be much more likely to know and care about (4*). However, this is not to say that we would ‘intuitively’ accept (4*) and reject (4**). In fact, if someone knowledgeable about cottontails and mosquitoes and virus transmission informed us about the truth of (4**), we would unhesitatingly and rightly accept it.

Sterken, however, insists that both (4*) and (4**) are false. Why? A central argument of hers is that we can generate contradictions by adding the clause ‘but typically/generally/normally they don’t’ to such weak generics, along the following lines:

Footnote 5 continued

My objection is valid independently of whether the condition is thought of in semantic or “worldly” terms. (I actually suspect that Leslie’s distinction is less clear than she assumes, but this is not an issue I can discuss here.)

⁶ To clearly see why (4**) would be just as true as (4*), imagine that a medical entomologist presented her evidence for (4**), showing in detail what constitutes the *Culex tarsalis*’s XYZ-virus-carrying mechanism and how the transmission of the virus from mosquito to cottontail occurs. The following objection would be patently misplaced: ‘The evidence you have presented is convincing: the mechanism is there, the XYZ virus is indeed transmitted in the way you describe, and the stated frequency of individual virus hosts in the US *Culex tarsalis* population is plausible. However, there is still a huge gap in your argument: you also need to show that the XYZ virus is striking enough for (4**) to be true.’ Leslie’s theory involves a psychologizing of truth-conditions that distorts our familiar use of type B generics; it is as if entomologists and other scientists would have to seek cooperation with the psychology department as soon as they wanted to test hypotheses like (4*) and (4**).

(9) *Culex tarsalis* mosquitoes carry the West Nile virus, but typically they don't.

As Mahrad Almotahari has convincingly argued, this argument does not seem to work (Almotahari, forthcoming). Even if our knee-jerk reaction to (9) might be a vague sense of incoherence, this sense should be taken with a pinch of salt. For if (9) were a genuine contradiction, its inconsistency could not be eliminated by adding more information (adding q to p & $\text{not-}p$ does not make the resulting sentence consistent). However, we can fairly easily add information to (9) which makes our vague sense of inconsistency go away. Consider:

(9*) In the US, *Culex tarsalis* mosquitoes carry the West Nile virus, but since most of them inhabit circumstances that fail to trigger their disposition to do so, they typically don't.

Or, suppose a Mississippian tries to soothe her Swedish visitor's anxiety by saying,

(9**) Our *Culex tarsalis* mosquitoes aren't as dangerous as you might think, but you're right: they still carry the West Nile virus. Typically, though, they don't.

(9*) and (9**) seem perfectly fine (cf. Almotahari forthcoming, 2–3).

I want to focus on another aspect of Sterken's discussion, however — but before that, let me just notice that there is an interesting difference between the following two generics:

(4*) In the US, *Culex tarsalis* mosquitoes carry the West Nile virus.

(10) Acorns develop into oak-trees.

These are both true, and both weak: only a small minority of acorns actually develop into oak-trees. However, with regard to (10), there is an important reading according to which adding 'but they typically don't' really generates a contradiction. So, consider:

(10*) Acorns develop into oak-trees, but typically they don't.

Whereas there is nothing the matter with a *Culex tarsalis* mosquito that does not carry the West Nile, there is something the matter when an acorn does not develop into an oak-tree: it *fails* to develop into an oak-tree. Either there is something the matter with the acorn itself, or with the circumstances in which it is situated. Thus, if 'typically' means something like 'if nothing interferes with their proper development *qua* acorns', then (10*) is a contradiction. By contrast, if 'typically' just means 'mostly'—as Sterken seems to be presuming, and which is presumably the most natural reading in most contexts—then (10*) is no more contradictory than '*Culex tarsalis* mosquitoes carry the West Nile virus, but typically they don't'.

However, the point I really want to make with regard to Sterken's discussion is that her choice of examples makes her error-theory seem more plausible than it actually is. 'Sharks attack bathers', 'Rottweilers maul children', and 'Australian snakes are poisonous' really *are* quite doubtful, and far from 'intuitively' true. 'Sharks attack bathers' is similar to 'Birds fly' since it fails to discriminate between species of shark; but of course it's much worse than 'Birds fly', precisely because there are very few species of shark that are dangerous to human beings. 'Australian snakes are poisonous' is similarly indiscriminate: among the 170 species of snake that live in Australia, only

about a dozen are poisonous. The sentence is not straightforwardly false, but it is certainly a gross oversimplification. In a truth-seeking context, replacing it by ‘Some Australian snakes are poisonous’ is definitely recommendable: nothing is lost and much clarity is gained.

What about ‘Rottweilers maul children’? Here, a determinate breed of dog is discussed, so the problem is not a lack of specificity as in ‘Birds fly’ or ‘Sharks attack bathers’. And Sterken rightly notices that a dialogue such as the following involves genuine disagreement:

A: Rottweilers maul children.

B: Don’t be silly. There have only been a few isolated incidents.

However, as Sterken interprets B’s objection, it involves a purely quantitative claim: attacks have been very few. It is noteworthy, however, that B characterizes these attacks as *isolated incidents*. B is objecting against the claim that these incidents form a non-accidental pattern, and that it somehow belongs to the nature of Rottweilers *qua* Rottweilers to maul children. It is not at all far-fetched to interpret A’s generic statement as involving precisely such a claim. And it is quite possible to have a sensible debate about such things. Recently, Norwegian researchers investigated attacks by wild wolves on human beings (Linnell et al., 2021). They studied 491 attacks registered globally between 2002 and 2020. In 382 cases (of which 14 was mortal), the wolves were infected by rabies. In 67 cases (of which 9 were mortal), the wolves seem to have perceived the human being as a rival or a natural prey. In 42 cases (of which 3 were mortal) the wolves had been provoked or threatened by the human being. Among all these attacks, only 11 took place in North America and Europe, and only 2 people died (both in North America). “Considering that there are close to 60,000 wolves in North America and 15,000 in Europe, all sharing space with hundreds of millions of people it is apparent that the risks associated with a wolf attack is above zero, but far too low to calculate”, the researchers conclude. These numbers seem to give little support to any sweeping generic claim to the effect that wolves kill human beings, even if a more nuanced version of some such claim may well be fine and even true. By contrast, ‘Tigers kill people’, and even ‘Tigers eat people’, seem justified, even if they need some nuancing too—tigers are much more aggressive toward humans than wolves are.⁷

In any case, my overall objection against theorists such as Sterken and Leslie is that the specific contents of the examples chosen matter a great deal more than they think. They treat their examples as mere instances of a vast and supposedly homogeneous class of statements, ‘weak generics’ (or, in Leslie’s case, ‘Type B generics’ or ‘striking property generics’ (cf. footnote 4)), and this is why it can seem reasonable to present a general theory intended to cover all instances within this supposedly homogeneous class. This one-size-fits-all approach involves fetishizing the true/false distinction: as soon as we question the prior assumption that all weak generics (or all Type B generics) have truth-conditions whose general structure can be identified without looking very closely at differences between individual cases, the aspiration to provide a general theory—be it an error-theory, or a theory of Leslie’s sort—seems like a non-starter.

⁷ A caveat: as I explain in Sect. 3, the case of Rottweilers is not exactly similar to the case of wolves and tigers. It matters that Rottweiler is a *breed* of dog rather than a *species*.

3 Social generics

Let's return briefly to the Rottweiler example. Rottweiler is a *breed* of dog, not a species of its own. This is important, for it affects the sort of normativity involved in generic statements about Rottweilers. Throughout centuries, Rottweilers have been molded by human breeders so as to live up to certain norms. Such normativity is taken to the extreme in the standards formulated by kennel clubs. The American Kennel Club description of the temperament of Rottweilers is as follows:

The Rottweiler is basically a calm, confident and courageous dog with a self-assured aloofness that does not lend itself to immediate and indiscriminate friendships. A Rottweiler is self-confident and responds quietly and with wait-and-see attitude to influences in his environment. He has an inherent desire to protect home and family, and is an intelligent dog of extreme hardness and adaptability with a strong willingness to work, making him especially suited as a companion, guardian and general all-purpose dog. (American Kennel Club, 1990)

This official standard claims to be “a description of the ideal Rottweiler” (ibid.) So, the conflict between the kennel club description and Sterken's example, ‘Rottweilers maul children’ is as follows. The kennel club description does not exclude that individual Rottweilers may maul children, but it locates the source of such behavior in failures of breeding and/or training. If someone claims that Rottweilers maul children, a Rottweiler owner may respond by saying, ‘No, a well-bred and appropriately trained Rottweiler does *not* maul children’. Contrast this with how one might want to argue against the statement ‘Wolves attack humans’. Here the issue is not a matter of breeding and training. Rather, a wolf-defender will argue that unprovoked wolf attacks are due to interferences with the wolf's natural development or environment: her claim is that a healthy wolf who is not provoked and hasn't been deprived of her natural preys will not attack humans. By contrast, the Rottweiler owner will claim that human intervention is required for a Rottweiler to flourish *qua* Rottweiler; nature and nurture are inseparable in this case.

It is interesting to note that the normativity involved in generic claims about Rottweilers (and about other breeds of dogs, as well as about other varieties of domestic animals) can often be meaningfully debated. For example, it is quite arguable that norms issued by kennel clubs are corrupt, in that they treat dogs as if they were artefacts or objects of decoration, rather than living beings and companions.⁸ After all, a flourishing Rottweiler is not necessarily the same as a Rottweiler which scores high

⁸ Consider some further passages from the American Kennel Club's Rottweiler standards: The head should be “[o]f medium length, broad between the ears; forehead line seen in profile is moderately arched; zygomatic arch and stop well developed with strong broad upper and lower jaws. The desired ratio of backskull to muzzle is 3 to 2. Forehead is preferred dry, however some wrinkling may occur when the dog is alert”. The dog's color should be “[a]lways black with rust to mahogany markings. The demarcation between black and rust is to be clearly defined. The markings should be located as follows: a spot over each eye; on cheeks; as a strip around each side of the muzzle, but not on the bridge of the nose; on throat; triangular marks on both side of the prosternum; on forelegs from carpus downward to the toes; on inside of rear legs showing down the front of the stifle and broadening out to front of rear legs from hock to toes, but not completely eliminating black from rear of pasterns; under tail; black penciling on toes”. “The behavior of

at dog exhibitions. Rather, the argument goes, we have to look much more broadly and deeply at what it is for a Rottweiler to live a good life. Even if such a good life is certainly a life together with and shaped by human beings, it is misguided to try to capture it in the sort of prescriptions favored by kennel clubs. In line with this point, it is also arguable that breeding sometimes goes too far for the dog's own good, and that crossbreeding is therefore something to be actively pursued, against the wish of the kennels.

I cannot here go into details about the logic of such debates, but it seems clear that we are now entering a space in which it would be hopeless to make any very sharp or principled separation between fact and value, or between truth and human interests. Indeed, it is a sphere in which it is extremely difficult to pin down the exact function of the distinction between truth and falsity. One thing seems clear, however: other terms of adjudication will have equally pertinent roles, including the distinctions between reasonable and unreasonable, sensitive and insensitive, searching and superficial, wise and unwise—just to mention a few. If we want to understand what is at stake in discussions of the sort I have indicated, it is crucial to take into consideration this whole repertoire of distinctions and not take it for granted that the true/false dichotomy has a privileged significance and straightforward application.

Since Rottweiler is a breed of domestic dog and is as such tied to the human social world, 'Rottweilers maul children' can reasonably be called a *social generic*. Indeed, its social dimension is largely what makes for the special complications that I have just been gesturing at. Consider now another kind of social generic, namely, generic statements about human artefacts. Of course, artefacts are of many sorts, and I cannot here engage in any extensive inventory. I will briefly discuss a couple of examples. First:

(11) Electric bass guitars have strings.

This is a fairly straightforward generic. It is true, even if not all electric bass guitars have strings (the strings may have been removed for purposes of reparation, for example). Indeed, (11) is not only straightforwardly true, it also captures something essential about electric bass guitars. There is nothing mysterious involved in such talk of essence—the point is just that an electric bass guitar is a string instrument, and if it is deprived of its strings it is completely useless for the relatively specific purpose for which it is designed. By contrast, consider.

(12) Electric bass guitars are four-stringed.

Uttered in the mid-1950s, this generic would have been true. However, Fender introduced their first six-stringed model already in 1961, and in the mid-1970s five-stringed basses became popular and are now more or less standard in many musical styles. So, (12) is no longer true, and its truth was never a matter of essence: the introduction of five- and six-stringed bass guitars immediately made perfectly good sense as complements to the four-stringed variety. Nonetheless, debates over the appropriate number of strings are occasionally stirred up. Thus, I have heard it been argued with some

Footnote 8 continued

the Rottweiler in show ring should be controlled, willing and adaptable, trained to submit to examination of mouth, testicles, etc." (ibid.). And so on and so forth.

fervor that if James Jamerson and Jaco Pastorius, the two most legendary and influential bassists of them all, didn't need more than four strings, then surely five- and six-stringed basses are aberrations. Conversely, the bassist Anthony Jackson has long been a spokesman for the view that only with the six-stringed variety did the electric bass guitar come to full fruition—it is, after all, a *guitar* and not a miniature double bass. (As happy in-betweens, five-stringed basses seem to need no such staunchly essentialist defenders; their users are just pragmatically pleased about the extra depth provided by the added B-string.)

Obviously, such quarrels are somewhat puerile, and often the arguments are made tongue in cheek. Truth isn't really at stake; rather, there is a clash of personal ideals and preferences. Not that such personal ideals and preferences are unimportant—in this case, they hang together with what one takes to be one's musical role models and with one's sense of belonging to a certain musical style or tradition. Such things matter, but the occasional controversies over what number of strings an electric bass should *really* have are nonetheless quite harmless. And among bass players in general, an opposite attitude of cheerful tolerance is much more common: the more varieties of gear there are, the merrier.

4 Pernicious social generics

As we know, however, there are many examples of social generics which are far from harmless. I have been discussing what may seem like comparatively trivial instances having to do with Rottweilers and electric bass guitars, whereas the contemporary debate about the potential harmfulness of generics has been centered on social generics about human beings—and, in particular, on social generics that seem to express and promote essentializing stereotypes about certain kinds or groups of human beings. I will now say something about how the points I have made earlier matter to how we should handle such truly dangerous cases. Perhaps unsurprisingly, the overall gist of my suggestions will be that fetishizing the true/false distinction is detrimental to clarity about the conflicts and dangers at issue. I will argue that we need to work with a richer repertoire of concepts and distinctions in order to adequately capture what is at stake, and also realize that the issue of truth versus falsity is sometimes not very central at all. There are even cases in which we should say of those who promote certain generic stereotypes that they do not take truth seriously enough for their generic statements to qualify even as *false*.

The dialectic of my discussion in this section will be somewhat complex. First, I will briefly revisit and further explore some of my earlier examples. Then I will introduce some pernicious ones. One overall point I want to make is that pernicious generics can be pernicious in various ways, depending on how, exactly, they are understood. My earlier examples will be used to identify different patterns of thought that may resurface in the pernicious employment of generics about human beings. I will then proceed to make a second point, namely, that despite their differences, these patterns of thought all involve what might be called an *objectification* of human beings. More precisely, these patterns of thought all disregard the fact that human beings—unlike mosquitoes, flamingos, and Rottweilers—are self-interpreting animals whose self-understanding

cannot be neatly separated from who they are. By engaging in a critical discussion of Leslie's treatment of pernicious generics in Leslie, 2017, I will argue that both these points are crucial for an adequate conception of the particularly deep and complicated dangers involved in such generics.

Let's start, then, by considering the following thought-experiment. Imagine that the West Nile virus has been extinguished from the surface of the Earth. In such a case, 'US *Culex tarsalis* mosquitoes carry the West Nile virus' will no longer be true, since there is no longer any West Nile virus around for the mosquitoes to carry. What will still be true is that *Culex tarsalis* mosquitoes are disposed to carry the virus, were it to reoccur: the virus-carrying mechanism is still present in them, ready to do its work. However, we may also imagine that, due to some mutation, this virus-carrying mechanism is regressed, so that the *Culex tarsalis* loses its capacity to carry the virus. Let us further assume that none of these changes has any other substantive effects on the rest of the life of *Culex tarsalis* mosquitoes. They live in the same way and as long as before, their procreation remains at the same level, and none of their other central capacities or characteristics are affected.

Now there is a sense of 'nature' or 'essence', according to which it would be correct to say: *Culex tarsalis* mosquitoes that live in such a virus-free world, or whose virus-carrying mechanism has been regressed, have not lost anything essential. Their nature, qua *Culex tarsalis* mosquitoes, remains the same as before. The virus-carrying was just an accidental appendage to their biological make-up.

Contrast this with the following case: World-wide pollution has made it the case that a majority of flamingos are born with a birth defect that makes it impossible for them to fly. The very use of the term 'defect' here shows that the case is different from the mosquito example: these flamingos are no longer able to flourish qua flamingos. What they have lost is not just an accidental appendage to their biological make-up, but something which is in an important sense essential to them qua flamingos (which is not to say that they are no longer flamingos, but only that they are defective ones). In such a world, the generic 'Flamingo fly' would not be false, but a qualification may be in place: Yes, flamingos fly, but nowadays more than half of them are incapacitated due to world-wide pollution.

Consider a third case, different from both the previous ones. A friend of Rottweilers reacts against the claim that Rottweilers maul children, by saying: 'No, even if there have been a few isolated incidents, Rottweilers are basically calm and confident dogs'. This Rottweiler-friend might well admit that if Rottweiler upbringing deteriorated—if Rottweiler owners stopped disciplining their dogs and let them grow up without any active training—it would indeed happen much more often that Rottweilers mauled children. Consequently, the dispute between the Rottweiler-friend and the Rottweiler-foe is not simply a dispute over what dispositions or mechanisms—inborn or induced by training—Rottweilers in fact have. Indeed, the Rottweiler-friend may agree with the Rottweiler-foe that Rottweilers have certain inborn dispositions toward aggressive behavior—dispositions that need to be disciplined. And the Rottweiler-foe agrees with the Rottweiler-friend that as things currently are, the aggressive dispositions of the vast majority of Rottweilers are curbed by training. What the Rottweiler-friend claims is that such discipline does not distort or curb the dog's nature, but, on the contrary,

brings it forth and makes the Rottweiler flourish *qua* Rottweiler. Unlike in the flamingo case, human intervention is *needed* for such flourishing to happen.

Now let's look at some pernicious generics. I will use some pretty vexing examples that figure frequently in the debate. Here is one:

(13) Women are submissive.

This statement can be understood in different ways. For example, it might be proposed as a claim about an accidental disposition among women. It would in this respect be similar to 'US *Culex tarsalis* mosquitoes carry the West Nile virus' as I construed it above, even if the disposition in question would not be thought of in biological terms; perhaps the idea is that the way women are brought up in today's society makes them submissive. Thus understood, the claim may well be made by someone who deplores the alleged fact that women are submissive, and who thinks that we should change our ways of bringing up women. Or, it may be made by someone who thinks it is a good thing that women are brought up in such a way that they become submissive. However, in neither case as I am imagining them is the alleged submissiveness tied to any idea of female flourishing.

Another possibility is that (13) is proposed as a claim about biological essence, similar to 'Flamingo fly'. 'Woman' is then treated as a biological category, and the claim is that individuals belonging to this category flourish by being submissive. Someone may make such a claim even while admitting that many or even most individual women today are not very submissive at all. He would then perhaps conceive of this contemporary situation as similar to the situation imagined above in which most flamingos have lost their capacity to fly due to world-wide pollution (even if he would presumably think of the 'pollution' which stops women from being submissive in cultural rather than biological terms).

A third possibility is that (13) is proposed as a claim similar to 'Rottweilers are calm and confident'. If so, the idea is not that women are submissive just in virtue of their biological make-up, but rather that they will develop such submissiveness if they are brought up in the right kind of way, and that this will be part of their flourishing *qua* women. Someone may make such a claim even while admitting that many or even most individual women are not very submissive at all, arguing that this only shows that the upbringing of women in contemporary society is in deplorable shape.

As galling as these different varieties of (13) are, its potential perniciousness has not yet been brought out in full. For, as I mentioned above, there is a feature of (13) which makes it different from talk about mosquitoes or flamingoes or Rottweilers, namely, the fact that it deals with human beings. And human beings are self-interpreting animals: What we are cannot always be neatly separated from what we take ourselves to be or what we aspire to become (Taylor, 1971). Thus, (13) can be used, not only to talk *about* how women should be brought up, but as itself an important *instrument* in such processes of upbringing—an instrument used to shape girls' self-understanding so that they adapt to and themselves come to embrace the stereotype expressed by the generic. By telling them that they are submissive, girls can be made to think of themselves as

submissive, and thereby become submissive—which, in turn, lends further credence to the original generic (this is what Ian Hacking famously calls a ‘looping effect’).⁹

Sometimes pernicious generics are used mainly as such instruments of upbringing. Consider:

(14) Boys don’t cry.

I presume no one would say that this generic is true of boys *qua* merely biological creatures. After all, it is characteristic of specimen of *homo sapiens* that they *do* cry, regardless of their sex (the natural-historical judgment ‘*Homines sapientes cry*’ is true). A more plausible reading aligns (14) with the ‘Rottweilers are calm and confident’ case, the idea being that only boys who have been brought up in ways that inhibit their biological disposition to cry will flourish *qua* men. However, even this may sometimes be an overblown interpretation—there are presumably cases in which (14) is not used to state any determinate ideas at all about what makes a boy flourish but is employed only to express and impose a certain preference with regard to what boys and men should be like. In such cases, the one who uses (14) comes closer to someone who stubbornly insists that electric bass guitars are four-stringed, but with the difference that the central function of (14) is to shape boys in accordance with the preference expressed and make them embrace such an ideal themselves. In this sort of case, the question whether (14) is *true* seems otiose. Its function is only to impose upon individual boys a certain contemptuous attitude towards boys who cry (including an attitude of self-contempt if they themselves cry) and thereby suppress their inborn disposition to do so.

Indeed, when we want to criticize the use of pernicious generics, it is important to keep in mind the possibility that *nothing* is being said which is sufficiently determinate to be either true or false. Suppose someone claims that women are submissive, whereupon we ask him to clarify what he means. Is he talking about some sort of widespread but accidental disposition, or does he have some concept of female flourishing in mind? Does he think of ‘women’ as a purely biological category, or as a category that needs to be understood in other, richer terms? In response, he just shrugs his shoulders, and has nothing clarificatory to say. In such a case, his claim is muddled rather than false. Which is not to deny its perniciousness and dangerousness.

Consider now,

(15) Muslims are terrorists.

This is one of Leslie’s central examples in her discussion of people’s tendency to generalize from very few striking instances to sweeping and highly dubious claims which have a generic shape. Now, Leslie and other participants in the debate agree that (15) is false. However, since (15) is a generic and not a universal generalization, apparently there is something to get nervous about here—for how are we to justify the claim that (15) is false, if we allow that other striking property generics such as ‘*Culex tarsalis* mosquitoes carry the West Nile virus’ are true? Of course, we may handle the worry by adopting Sterken’s error theory and argue that all these generics

⁹ Hacking 1999, p. 34; 2002, p. 106. Many theorists would talk about the ‘performative’ function of generics here. In a paper inspired by Austin, I prefer to avoid this widespread use of the term ‘performative’ which is substantively different from Austin’s own (in Austin 1975).

are false. However, the more common strategy is to develop some general theory of what truth-conditions such generics have and use that general theory to explain why ‘Muslims are terrorists’ and other pernicious ones are false.

Leslie proposes such a strategy in Leslie, 2017. According to her analysis, striking property generics such as ‘Muslims are terrorists’ and ‘*Culex tarsalis* mosquitoes carry the West Nile virus’ involve the claim that the kind in question—*Muslim* in the first case, *Culex tarsalis mosquitoes* in the second—is a good predictor of the ascribed property. Thus, ‘Muslims are terrorists’ involve the claim that Muslims, *qua* Muslims, are disposed to be terrorists—the idea being that even those Muslims who do not in fact engage in terrorist activities are nonetheless liable to do so given that the right circumstances are present. Similarly, ‘*Culex tarsalis* mosquitoes carry the West Nile virus’ involves the claim that *Culex tarsalis* mosquitoes, *qua Culex tarsalis* mosquitoes, are disposed to carry the West Nile—the idea being that even those *Culex tarsalis* mosquitoes that do not in fact carry the virus will do so if exposed to the virus in the relevant fashion. So, “A generic statement in which a striking property is predicated is, I claim, true if and only if some members of the kind in question possess the relevant property, and the others are typically disposed to possess it” (Leslie, 2017, p. 404). According to Leslie, this truth-conditional analysis gives us the resources to make clear why ‘*Culex tarsalis* mosquitoes carry the West Nile virus’ is true whereas ‘Muslims are terrorists’ is false.¹⁰

Leslie’s proposal may seem unobjectionable. Her notions of ‘good predictor’ and ‘disposition’ are obviously meant to be construed broadly, as covering cases of very different sorts (“biological”, “cultural”). Conceived abstractly enough, it may seem undisputable that her analysis captures a general pattern present in all those instances of generic thought that she is talking about.

However, I want to raise two worries. First, let me repeat that we must not take it for granted that a generic, just because it has the shape of a declarative sentence, is either true or false. Indeed, a spontaneous reaction to Leslie’s discussion of (15) is to ask: How many people would actually say that ‘Muslims are terrorists’, just as it stands, is straightforwardly, unqualifiedly *true*? And how many would base their claim *only* on having heard about one or very few terrorist attacks by Muslims? I do not doubt that *some* would. However, in such cases their reasoning is arguably so muddled and careless that the best thing to say is not that their claim is false, but that their minds are clouded to such an extent that we must question their capacity to speak clearly enough about these things for the terms ‘true’ and ‘false’ to be applicable at all (even if their capacity for precision might work well when topics other than Muslims are at issue). If they really have *nothing* more to say—if they provide no support of it—dignifying their claim even with the epithet ‘false’ means seriously downplaying the amount of confusion involved.

My second worry is that the general character of Leslie’s analysis constitutes a more serious problem than what appears at first sight. For the differences that it abstracts from go deeper than Leslie can afford to admit. This will require some explaining.

¹⁰ For simplicity’s sake, I drop the reference to specifically North American *Culex tarsalis* mosquitoes in what follows. None of this makes any relevant difference to the points I will make.

Above, I noticed a difference between ‘*Culex tarsalis* mosquitoes carry the West Nile virus’ and ‘Flamingos fly’: whereas the virus-carrying mechanism is just an accidental appendage to the biological make-up of *Culex tarsalis* mosquitoes, the capacity to fly is part of what makes flamingos flourish *qua* flamingos. In this sense, flying belongs to the nature or essence of flamingos, whereas carrying the West Nile does not belong to the nature or essence of *Culex tarsalis* mosquitoes. This sort of difference is made invisible in Leslie’s account. She does say that generic claims of the sort she is talking about involves essentialization, but the concept of essence that she is working with is tied to a notion of causal grounding rather than to any notion of flourishing. According to her account, proponents of a striking property generic take the relevant disposition to be part of a hidden nature shared by members of the kind—a nature which causally grounds the outwardly observable behavior of these members. Consequently, given Leslie’s conception of essence, the virus-carrying mechanism is no less essential to *Culex tarsalis* mosquitoes than the capacity for flight is to flamingos.¹¹

This, in turn, makes Leslie blind to the fact that in many cases, controversies over a generic ‘As are F’ are not just controversies over what dispositions are present in some already delineated class of entities (the As). Instead, many such conflicts involve disagreements over what makes As flourish *qua* As—what constitutes an exemplary or prototypical A. Thus, consider again our imagined Rottweiler-friend who insists that Rottweilers do not maul children but are calm and confident. As I said earlier, this Rottweiler-friend might well agree (i) that there have been cases in which children have been mauled by Rottweilers, and (ii) that Rottweilers have an inborn disposition to aggressivity which, if it isn’t properly disciplined, might be turned against children. So, if Leslie is right, it would seem as if such a Rottweiler-friend has thereby admitted that Rottweilers maul children—for doesn’t admitting the truth of (i) and (ii) amount to admitting that “some members of the kind in question possess the relevant property, and the others are typically disposed to possess it”?

It is clear, however, that the generic statement made by our Rottweiler-friend requires a different reading. Again, what he is saying is that appropriately trained Rottweilers—ones trained in such a way that they flourish *qua* Rottweilers—do not maul children. Now, it may still seem as if this isn’t a problem for Leslie, for she might respond that what our Rottweiler-friend is doing here is just to seek a more restrictive generalization which cites a better predictor for mauling/not mauling children: *properly trained Rottweilers* do not maul children, whereas *poorly trained Rottweilers* do (cf. Leslie, 2017, p. 413). The problem with this response, however, is that if we ask our Rottweiler-friend to specify what makes a certain training-program for Rottweilers appropriate, he will say that one necessary condition for its appropriateness is precisely that it produces Rottweilers that do not maul children. From the perspective of Leslie’s analysis, in which predictability is the central notion, such a Rottweiler-friend’s position will seem empty as it is based on what will appear like a question-begging identification of the predictor: If an individual Rottweiler mauls a

¹¹ Leslie does note that she is working with a notion of essence used by psychologists rather than by philosophers (2017, p. 406, n. 29). As I will argue below, however, she does not realize that her conception of essence obliterates differences that are highly relevant if we want to understand how pernicious generics should be handled.

child, it has *a fortiori* not been properly trained, so it follows trivially that properly trained Rottweilers do not maul children.

However, the Rottweiler-friend's claim isn't empty, for two reasons. First, he is saying that there *are* in fact ways of training Rottweilers that are appropriate. Rottweilers are not hopeless cases (as, presumably, tigers are)—they *can* be trained so that they do not maul children. Second—and this is a crucial point—our Rottweiler-friend is saying that these training-programs are appropriate not only because they produce Rottweilers that do not maul children, but because they produce Rottweilers that flourish in a much richer sense. After all, suppose our Rottweiler-friend tries to justify his claim that Rottweilers do not maul children by citing the fact that if we regularly give Rottweilers electric shock therapy, starting when they are very young, they will end up as docile zombies that never maul children. This would of course be ridiculous as evidence that Rottweilers do not maul children. Why? Because such a program of electric shocks is a patently inappropriate way of bringing up Rottweilers. Why? Because even if it results in Rottweilers that do not maul children, it also means that they are destroyed *qua* Rottweilers.

Now we can see very clearly why the debate between a Rottweiler-friend and a Rottweiler-foe over the generic 'Rottweilers maul children' will have as an absolutely central component the question of what makes a Rottweiler flourish *qua* Rottweiler. We can also see why a general analysis such as Leslie's, which focuses on mere predictability, hinders us from adequately understanding how the issue should be dealt with since it hides the whole dimension of flourishing from view. By taking a scheme which does fit the '*Culex tarsalis* mosquitoes carry the West Nile virus' and then impose it onto the 'Rottweiler maul children' case, her account makes invisible the significance that the notion of flourishing has in the latter. As I said above, questions about what makes a Rottweiler flourish can perfectly well be meaningfully debated, but in such debates the notions of truth and falsity will be inextricably interwoven with other terms of adjudication (reasonable/unreasonable, sensitive/insensitive, searching/superficial, wise/unwise, and so on) in ways that undercut the idea that a general truth-conditional theory can do justice to the complexities involved. And these are complexities of a sort that are crucial if our aim is to handle the dangers involved in the use of many pernicious social generics.

When it comes to social generics about human beings, further important complications arise, due to the fact that such generics shape people's self-understanding. As I noticed above, people who claim that properly brought up women are submissive and that properly brought up boys don't cry will presumably also think that it will be part of such proper upbringing that the women and boys are themselves made aware of and brought to embrace these generics. Someone who disagrees with the generics will regard this as a form of indoctrination, whereas the women or boys who have been successfully brought up in this fashion will themselves reject such a charge as patronizing. In this familiar sort of conundrum, assuming that 'Women are submissive' and 'Boys don't cry' can be rejected because they fail to fulfil truth-conditions specified in terms of predictive power is not just wrong-headed but positively detrimental to clarity and progress. For the issue here is one of understanding rather than prediction—of hermeneutics rather than science, if you like. Hence something much more difficult is required, namely, entering a dialogue with those whose self-understanding one regards

as unwise or muddled—a dialogue meant to invite one’s interlocutors to see for themselves that there are other and better notions of flourishing available. The point need not be to replace one uniform conception of female or male flourishing with another, but rather to exhibit and clarify an open-ended plurality of possibilities. In such a dialogue, the rendering of particular examples, stories, life experiences, and so forth, will be crucial, not just because they are useful pedagogical tools, but because the disagreement itself *is* a disagreement over the significance of such examples, stories and experiences in human life.¹²

Let me now finally return to the example ‘Muslims are terrorists’. I noted above that one might wonder how common it actually is that people make such a claim *only* on the basis of having heard about a few instances of terrorism carried out by Muslims. I do not doubt that this happens, but arguably a more interesting and often more dangerous kind of case is people who agree that ‘Muslims are terrorists’ is indeed a hyperbolic, coarse and perhaps even Islamophobic claim, only to go on to say that it nonetheless contains a grain of truth. Asked to identify this grain, they might give different answers. They may try to explain their point in purely statistical terms: People who grow up in predominantly Muslim countries have a greater likelihood of becoming terrorists than people who grow up elsewhere. If so, their claim is no longer a generic, and needs to be scrutinized as the statistical claim that it is. More commonly, however, what happens is that the alleged grain of truth is spelled out partly in terms of what constitutes an exemplary, prototypical Muslim—a Muslim who, as it were, flourishes *qua* Muslim, and in this sense embodies Islamic religion and culture. Do *such* Muslims endorse the use of violence against non-Muslims, and perhaps even against other Muslims whom they regard as apostates? And if so, shouldn’t we worry that their religion and culture can play a destructive role, especially in times when there are also many other tensions—economic, political—between the Muslim and non-Muslim parts of the world?

As they stand, these questions are of course still quite crude. Moreover, when it comes to the question of who is an exemplary Muslim, Muslims must themselves be allowed to make their voices heard in the discussion—for the notion of an exemplary Muslim, a Muslim who flourishes *qua* Muslim, cannot be separated from how Muslims themselves understand what it is to be a Muslim. Thus, it is no coincidence that people who make ambitious attempts to identify the alleged grain of truth in the pernicious generic give detailed descriptions of contemporary Muslim societies, provide learned accounts of Islam’s history, discuss various passages from the Quran, make long interviews with Muslim scholars, and so on and so forth.¹³ However, given the extremely multifaceted character of the world-wide Muslim community and its sprawling history and tradition, and given the fact that the question ‘Who is an exemplary Muslim?’ is

¹² The significance and difficulties of conundrums of this sort are familiar from many works in hermeneutical and Marxist traditions of thought, as well as from key works in feminist philosophy. For two classical, lively and illuminating discussions, see de Beauvoir 1976 (especially part III) and Taylor 1971 (especially pp. 15 ff.). Let me emphasize that a dialogue of the sort I am sketching is often not sufficient; in many cases, such a dialogue can be meaningfully pursued only if there is also a liberation from material means of oppression.

¹³ For a both thought-provoking and disturbing example, see Wood 2015. Wood challenges Barack Obama’s claim that the Islamic State is “not Islamic”, arguing in considerable detail for the opposite view: “The reality is that the Islamic State is Islamic. Very Islamic.” Thanks to an anonymous referee for the reference.

itself a contentious issue *within* this community and tradition, the situation becomes extremely messy. There is also a worry about the question-begging nature of the disagreements involved—for when it comes to who is an exemplary Muslim, shouldn't the views of *exemplary* Muslims have more weight than the views of non-exemplary Muslims? I am not saying that these messy discussions are necessarily confused, even if they often are. My point is rather that it is difficult to find any foothold here for the application of an analysis such as Leslie's, or for any other general truth-conditional theory of generics. Indeed, the question how and when the true/false distinction applies at all in this area is itself a contentious issue within these conflicts, having to do with the question to what extent Islam should be seen as a creed with one correct interpretation or whether it is better conceived as a socially and historically malleable tradition whose significance depends on what its believers make of it. Getting anywhere here will often be a matter of seeking understanding in a dialogical, piecemeal fashion, rather than to impose some allegedly neutral truth-condition such as predictability.

5 Semantic modelling versus piecemeal attention to details

Alfred Ayer once accused Austin of acting like a greyhound that doesn't like running and therefore bites all the other dogs so that they cannot run either (Berlin, 2014, p. 153). Perhaps my discussion provokes a similar reaction. I may seem to want to complicate matters so much that any attempt to say anything of significance in the area of generics looks futile. However, this has not been my aim. Rather, what I have been arguing is that theorists are often running in the wrong direction. As Jennifer Saul says in one of the most thoughtful papers published recently within this field, “we need to get better at talking and thinking about [generics]” (Saul, 2017, p. 14). According to Saul, this means that “[w]e need to press people to spell out their evidence for their generic claims and to reflect on what that evidence really means” (ibid.). I agree but want to go further. Getting better at talking and thinking about generics also means getting better at recognizing their heterogeneity. This is possible only if we reject the prejudice that our main aim should be to find truth-conditions whose general structure can be identified without looking very closely at differences between individual cases. Such a focus may seem required by the supposed need to come up with a formal semantics for generic statements. If I am right, however, a formal approach of this sort often blocks the path towards the kind of rich and nuanced understanding we need in order to grasp the significance of generic statements in real life and get clear about how they can be dangerous. Let me end this paper by saying a bit more about this.

To begin with, I don't mean to suggest that contemporary theorists show no appreciation at all of the fact that the significance and danger of generics goes beyond what a standard truth-conditional semantics can be expected to handle. For example, Sterken has developed a contextualist account of generics which is more general than her specific view on weak generics, and more open to heterogeneities of usage. Recently, she has merged this contextualist account with David Plunkett's and Timothy Sundell's conception of metalinguistic negotiation to create a framework which allows us to think in complex ways about the significance and dangers of generics in various real-life contexts (Plunkett et al., 2023). Another example of such a richer approach

is Sally Haslanger's in many ways illuminating discussion of how ideology critique does not consist in a mere rejection of false statements but is transformative in a more radical sense—a discussion in which social generics play a central role (Haslanger's central example is 'Crop-tops are cute!') (Haslanger, 2007).

My sense, however, is that Sterken and Haslanger still allow the true/false distinction to play a structurally fundamental role that is questionable. Thus, the variability introduced by Sterken's contextualism remains standardly truth-conditional: Even if contextual parameters are taken to fix what a generic sentence expresses on a given occasion of use, this expressed content is still accounted for in terms of the conditions under which the utterance is true or false. Adding metalinguistic negotiation to the mix does not challenge this assumption, for what such negotiations are about, on this model, is what the uttered generics should be used to express—where, again, what is expressed gets cashed out in terms of dichotomously conceived truth-conditions. With regard to Haslanger, even if she rightly points out that a clash between ideologies cannot be understood simply as disagreement over whether certain statements are true or false, she nonetheless seems to take it for granted that before such a clash happens—when an ideology or social practice is not yet criticized but viewed as it were parochially, from within the practice or the social milieu—the true/false dichotomy does capture how the parochial users themselves understand the statements in question. By contrast, I have argued that the true/false dichotomy often fails to capture what is going on even within the object language or the social milieu. The distortion starts at an even more basic place than Sterken or Haslanger assume, and thus they remain under the spell of the true/false fetish even if they try to resist it.¹⁴

However, it may still be argued that Ayer had an important point against Austin, and against me. For why expect that semantic theories should capture language in all its marvelous complexity? Isn't it the other way around: Precisely because language is so incredibly complex, the only reasonable and fruitful approach in semantics is to construct partial models that involve a great deal of abstraction and idealization. Like in any other special-science modelling, distortions and exceptions are bound to arise—but as long as we are methodologically clear about what we are doing, wherein lies the harm?¹⁵

This is a powerful objection, not least because it would seem to involve a sophisticated appropriation of my own skepticism towards the true/false fetish, albeit now transformed to apply to the status of semantics itself. As Gabe Dupre explains, models

provide an indirect way of investigating the world. By producing a model, we create a surrogate for its real-life target. We can then study the properties of the

¹⁴ This point opens up the possibility of a criticism of the sort of relativist model which Haslanger ends up defending towards the end of her 2007 paper. As I think of it, there is a *continuity* between 'parochial', milieu-internal usage and ideology critique which undermines such relativism. Consider again my point about how kennel club norms can be meaningfully debated. This is not a debate at a meta-level but illustrates how a sensible ideology critique can arise *from within* a social practice or milieu, as it makes use of a wide repertoire of terms, including 'reasonable', 'sensitive', 'superficial', and 'wise' – terms that the participants can be taken to share to an extent which makes a relativist construal of the debate artificial. Haslanger's notion that such debates can be resolved by "forming or finding a common milieu" (Haslanger 2007, p. 87) hides the common ground *which is already there*. I plan to develop this criticism of Haslanger in a separate paper.

¹⁵ Thanks to an anonymous referee for pressing me to address this issue.

model and, under certain conditions, make inferences about the properties of the target. Of course, no model will perfectly resemble the target. If it did, there would be no benefit to studying the model in the first place. (Dupre, 2020, p. 9; cf. also Yalcin, 2018)

As a consequence,

models call for a different mode of evaluation than theories. Theories may be true or false. However, such evaluations are unhelpful in discussing models. Because models invariably misrepresent, or at least only partially represent, their targets, they do not really aim at truth. Instead, models are best evaluated on the basis of how useful they are. (Dupre, 2020, pp. 9–10)

So, what if it is *I* who am being wedded to the true/false fetish, in taking it for granted that the aim of semantics is truthful representation rather than useful modelling? If we instead think of semantics in terms of modelling, then isn't my claim that formal truth-conditional semantics fails to capture all the nuances of our real-life use of generics trivial and beside the point?

This attempt to dismiss my criticisms strikes me as too quick. However, let me first say that I applaud the conception of semantics as model-based. Indeed, if criticism of the sort I have delivered in this paper can help induce such methodological clear-sightedness and modesty among philosophers and semanticists, that is perhaps no minor achievement—for such self-awareness might still not be as widespread as one could wish. Be that as it may, however, it also needs to be said that thinking of semantics in terms of modelling does not by itself suffice to fully accommodate the points I have made. For scientists may be aware that they are working with models and yet be insufficiently sensitive to their limitations. Such insensitivity might manifest itself in the fact that models of a certain sort are given an unduly dominant position in the field, whereas alternative approaches are marginalized. This sort of point is familiar from economics—just consider the multifarious criticisms of how formally neat neo-classical models are allowed to dominate at the expense of alternative viewpoints which are less neat but can provide richer and more nuanced accounts of economic behavior (including aspects that the conventional models have tended to distort or ignore, such as the limitation of the earth's natural resources and the existence of genuine altruism). Similarly, the Austinian type of criticism, of which my discussion in this paper is an instance, can be used to challenge the dominance of standard truth-conditional formal models in the contemporary study of language.¹⁶

A crucial issue in this connection is how the usefulness of models is to be conceived and assessed. One way of securing the dominance of certain models is to invoke a measure of usefulness which itself guarantees that those models are going to score higher than alternative approaches. For example, suppose we say that a useful semantic model is one which “for the purposes of the current inquiry, adequately reproduces the behavior of the natural language phenomenon it targets, and suggests as yet unobserved behavior of such a phenomenon which can then be tested” (Dupre forthcoming,

¹⁶ Of course, I have not myself even begun to develop any alternative model in this paper. At most, my critical discussion amounts to a preamble to such model-construction, indicating why it might be worth pursuing.

10; emphasis removed). So far, so good. But what, exactly, do we mean by “the behavior of the [targeted] natural language phenomenon”? Suppose it turns out (i) that this behavior is conceived in terms of what truth-conditional content gets expressed by the sentences or utterances that we are interested in, (ii) that the true/false distinction is construed dichotomously, and (iii) that the possible relevance of other terms of assessment—including ‘rough’, ‘exaggerated’, ‘oversimplified’, ‘good enough’, ‘superficial’, ‘wise’, ‘sensitive’, and so on and so forth—is not considered at all but just ignored (or perhaps brushed aside as of merely “pragmatic” significance). If so, we have stacked the cards in favor of standard truth-conditional semantics and made invisible the usefulness of approaches that are more responsive to the sort of points I have made in this paper, and which do not hide but instead bring out the heterogeneity of real-life language use.¹⁷

It may be responded that the very idea of incorporating terms of assessment such as those mentioned under (iii) above into semantic modelling would lead to chaos. Indeed, who can reasonably believe that such terms of assessment can be neatly captured by some model that lives up to formal standards of systematicity and rigor? My counterquestion would then be: Shouldn’t one rather conclude that such formal standards of systematicity and rigor are not always appropriate? Wouldn’t insisting on such standards across the board mean that our modelling runs a serious risk of deteriorating into an exercise whose immunity to external criticism is bought at the expense of genuine empirical significance?

Arguably, a better response is to allow and encourage the development of less formal models—models that incorporate a richer repertoire of terms of assessment than what standard truth-conditional semantics offer. After all, even if such models would not exhibit the same rigor and systematicity as the truth-conditional models we are used to, this does not necessarily mean that they are chaotic. In fact, science has seen many examples of models that do not live up to such strict formal standards but are still very fruitful (for an interesting discussion of some examples from biology, see Godfrey-Smith, 2006).

However, two things are important to keep in mind here. To begin with, what I have been arguing in this paper is not that terms of assessment such as ‘rough’, ‘good enough’, ‘sloppy’, ‘superficial’, ‘sensitive’, ‘wise’ and so on, should be considered simply *in addition* to ‘true’ and ‘false’ when we try to understand how generics function (and thus be treated as belonging to a separable sphere of “pragmatics”, say). I have made a deeper point, namely, that these other terms of assessment *interact* with our use of ‘true’ and ‘false’ in multifarious ways—ways that do not fit the simplistic uniform construal of the true/false distinction that characterize standard truth-conditional semantics. My claim is that the functions of ‘true’ and ‘false’ in our real-life use of generics can come into clear view only if we look at how they are related (not only

¹⁷ Dupre conceives the usefulness of semantic models in terms of (i)-(iii). However, in the specific dialectic context of his paper – in which he argues against forms of radical contextualism *that conceive of expressed content in similarly truth-conditional terms* – this might be in order. As I have already noted, my own approach differs from such radical contextualism by questioning the idea that the real-life role of the true/false distinction can be adequately understood in this way. My aim is to undermine standard truth-conditional models at a more fundamental level than what radical contextualists of the sort Dupre discusses can reach.

to each other but) to these other terms of assessment as well; and these relations are very complicated and look different in different cases. Consequently, my point is not simply that standard formal truth-conditional models need to be *supplemented* with other (presumably less formal) models that deal with a wider repertoire of assessments. Rather, what I am arguing is that as long as the true/false distinction is treated separately from this wider repertoire of assessments, its significance in our real-life use of generics will be distorted.

Again, one might respond that distortion is inevitable in semantic modelling, and that a certain division of labor is necessary if we want to get anywhere. This is fine, as far as it goes—it would indeed be foolish of me to try to issue detailed instructions for semantic modelling. Let me instead conclude by bringing in my second point, namely, that the idealization and simplification involved in such modelling is not always what we need. For example, when we try to get clear about the various ways in which social generics can be dangerous, a piecemeal, descriptive approach where we look patiently at individual examples and proceed from the understanding we already have as competent and sensitive language users, can be just as, or even more, illuminating. For the dangers are often in the details, and these details will come into clear view only if we appreciate how heterogeneous and complicated our use of generics actually is. In such cases, getting better at talking and thinking about generics means getting better at attending to precisely the sort of distinctions and differences that semantic models tend to ignore in the name of simplicity and unification. Engaging in such detailed attention is then a way of taking truth seriously rather than fetishizing it.

Acknowledgements An earlier version of this paper was presented at the conference “Philosophy’s Experimental Turn and the Challenge from Ordinary Language” at the Freie Universität Berlin. Thanks to the organizers, and to the participants for a valuable discussion. Earlier versions were also presented Åbo Akademi University, Stockholm University and Tel Aviv University; I thank the audiences for their comments. I especially thank Ulrika Björk, Maria Balaska, Cora Diamond and Camilla Kronqvist for their friendly criticisms of earlier drafts. Two anonymous referees for this journal gave extensive and very helpful comments on the first submitted version of the paper; I thank them for encouraging me to clarify and develop some of the central ideas in more detail.

Funding Open access funding provided by Åbo Akademi University (ABO).

Data availability No datasets were generated or analyzed during the current study.

Declarations

Conflict of interest The author has no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Almotahari, M. (Forthcoming). Weak Generics. Forthcoming in *Analysis*.
- American Kennel Club. (1990). Official Standard of the Rottweiler. images.akc.org/pdf/breeds/standards/Rottweiler
- Austin, J. L. (1975). *How to do things with words* (2nd ed.). Harvard UP.
- Berlin, I. (2014). *Personal impressions* (updated). Princeton UP.
- Cohen, A. (2012). Generics as modals. *Recherches Linguistiques De Vincennes*, 41, 63–82.
- Colpitts, T. M., et al. (2012). West Nile virus: Biology, transmission, and human infection. *Clinical Microbiology Reviews*, 25, 635–648.
- de Beauvoir, S. (1976). *The ethics of ambiguity*. Kensington Pub.
- Dupre, G. (2020). Idealisation in semantics: Truth-conditional semantics for radical contextualists. *Inquiry*, online first.
- Godfrey-Smith, P. (2006). The strategy of model-based science. *Biology and Philosophy*, 21, 725–740.
- Hacking, I. (1999). *The social construction of what?* Harvard UP.
- Hacking, I. (2002). *Historical ontology*. Harvard UP.
- Haslanger, S. (2007). ‘But mom, crop-tops are cute!’ Social knowledge, social structure and ideology critique. *Philosophical Issues*, 17, 70–91.
- Lasersohn, P. (1999). Pragmatic halos. *Language*, 75, 522–551.
- Leslie, S.-J. (2008). Generics: Cognition and acquisition. *Philosophical Review*, 117, 1–47.
- Leslie, S.-J. (2017). The original sin of cognition: Fear, prejudice and generalization. *Journal of Philosophy*, 114, 393–421.
- Linnell, J. D. C., Kovtun, E., & Rouart, I. (2021). *Wolf attacks on humans: an update for 2002–2020*. NINA report 1944. Norwegian Institute for Nature Research.
- Longino, H. E. (1994). In search of feminist epistemology. *The Monist*, 77, 472–485.
- Nickel, B. (2016). *Between logic and the world*. Oxford UP.
- Plunkett, D., Sterken, R. K., & Sundell, T. (2023). Generics and metalinguistic negotiation. *Synthese*, 201, 1–46.
- Saul, J. (2017). Are generics especially pernicious? *Inquiry*, online first.
- Sterken, R. K. (2015). Generics, content, and cognitive bias. *Analytic Philosophy*, 56, 75–93.
- Taylor, C. (1971). Interpretation and the sciences of man. *Review of Metaphysics*, 25, 3–51.
- Teichman, M. (2015). *Characterizing kinds: A semantic for generic sentences*. University of Chicago.
- Thakral, R. (2018). Generics and weak necessity. *Inquiry*, online first.
- Thompson, M. (2008). *Life and action*. Harvard UP.
- Wood, G. (2015). What ISIS really wants. *The Atlantic*, March 2015 issue. Retrieved at <https://www.theatlantic.com/magazine/archive/2015/03/what-isis-really-wants/384980/>
- Yalcin, S. (2018). Semantics as model-based science. In D. Ball & B. Rabern (Eds.), *The Science of meaning: Essays on the metatheory of natural language semantics* (pp. 334–360). Oxford UP.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.