Generic Cognition: A Neglected Source of Context Sensitivity

Forthcoming in *Mind & Language*

Mahrad Almotahari
University of Edinburgh
School of Philosophy, Psychology and Language Sciences
Department of Philosophy
40 George Square
EH8 9JX
malmotah@ed.ac.uk

No funding bodies contributed to this research.

Abstract

What is the relationship between the claim that generics articulate psychologically primitive generalizations and the claim that they exhibit a unique form of context sensitivity? This paper maintains that the two claims are compatible. It develops and defends an overlooked form of contextualism grounded in the idiosyncrasies of System 1 thought.

Generics Bare Plurals Default Generalizations Context System 1

This paper is about two ongoing research programs—one in cognitive science and one in philosophy of language. Both are concerned with the meaning of generic sentences, but they approach the subject from very different starting points. A central claim of the former is that generics articulate psychologically primitive or default generalizations.[1] Call this view *primitivism*. A key component of the latter is that generics exhibit a ubiquitous and unique sort of context sensitivity.[2] Call this view *contextualism*. The question that interests me is simple: what is the relationship between these two views?

Critics of primitivism maintain that they are largely incompatible—not in a strict formal sense, but in one that I will make precise very soon—and they argue at length in support of contextualism.[3] If they are right, then we are in the unfortunate position of having to reckon with the highly compelling evidence for primitivism. I know of one serious attempt to do so from a contextualist point of view. But I will argue that it fails. We would be much better off, then, if primitivism and contextualism were compatible.

[1] See Gelman (2003), Leslie (2007, 2008, 2012, 2015*ab*, 2017, 2022), and Leslie et al. (2011).

[2] See Nickel (2009; 2010; 2012; 2016), Sterken (2015*c*), Thakral (2018), Nguyen (2020), Hesni (2021), Lee and Nguyen (2021).

[3] See Sterken (2015*abc*), Nguyen (2020), and Lee and Nguyen (2021). For an expression of sympathy, see Saul (2017).

One of my central claims is that they are. To vindicate this claim, I will address a forceful argument for incompatibility.[4] The argument has never been explicitly formulated, as far as I know, but I will not be able to do so until I provide a more concrete understanding of what a default mode of generalization is. Eventually, I will explain why this argument, too, is mistaken. But that will take some time.

My project begins in Section 1 with a crucial preliminary question. What is a generic sentence? The issue is not as straightforward as one might think, and I suspect that critics of primitivism rely on a much broader view of the target phenomenon than is warranted. My conception of the subject matter is based on the conviction that "in scientific inquiry, we make progress by severely narrowing our focus" (Yalcin 2018, p. 351, paraphrasing Frege). Given the empirical nature of our investigation, the task of demarcating the relevant evidence proceeds in lockstep with substantive theorizing (Chomsky 1980*a*, p. 15). Section 2 provides a more detailed characterization of both primitivism and the larger theoretical framework to which it belongs. Section 3 sketches (what I take to be) the master argument for the framework. With it in view, I will be able to explain why a recent contextualist attempt to assimilate some of the framework's virtues fails. Section 4 highlights an aspect of the framework that its critics erroneously neglect, namely, that it predicts widespread context sensitivity of a unique sort, thus vindicating a form of primitivist contextualism. Section 5 concludes.

At every stage, my reliance on the groundbreaking work of Sarah-Jane Leslie will be obvious. By spotlighting the underappreciated aspects of her psychological framework, I hope to address many of the most influential objections levelled against it. The upshot is a form of contextualism grounded in the nature of generic cognition, rather than the semantics of *Gen*. The fundamental idea guiding this work is that our default generalizations are sensitive to psychological factors that make generics essentially contrastive. These factors (in particular, salience and stability) determine the contextually relevant contrast class with respect to which a generic generalization is interpreted and assessed.

1. The Target

What is a generic sentence? How should we define the target phenomenon? These are hard questions. Discussion often centers around some familiar examples:

(1) Tigers have stripes.
(2) Birds fly.
(3) Bees reproduce.
(4) Orange-Crusher 2000s crush oranges.
(5) Firefighters fight fires.
(6) Mosquitoes carry the West Nile virus.
(7) Pit bulls maul children.
(8) Cars have radios.
(9) Bottles have necks.

[4] The argument was raised, more or less in the form I present it, by Bernhard Nickel (pc).

(10) Supreme Court Justices have even Social Security numbers.

Each sentence contains a bare plural noun phrase in subject position. But generics come in many different shapes. Some contain a mass noun in subject position (*Candy rots teeth*). Some contain a singular definite (*The whale is a mammal*). And some contain an indefinite (*A madrigal is polyphonic*). Going forward, I will focus on bare plurals.

Semantically, (1)-(10) are somewhat heterogeneous. But the differences between them should not obscure an important similarity: each example instantiates a surface structure, *Ks are F*, that somehow expresses a generalization about Ks as such, not just a particular subset of them. Presumably, that is why they are called *generic* sentences—the label is supposed to make this aspect of their meaning (that they involve a genus- or category-wide claim) salient.[5] It is also why a pit bull enthusiast can challenge the truth of (7) by insisting that there are no bad pit bulls as such, only bad owners.

With these observations in mind, compare sentence (1) with *Tigers are on the front lawn*. Although the latter instantiates the same superficial form, it is not about tigers as such. Plausibly, it is not a generic sentence (Leslie 2008, p. 4). Rather, it is an existential in disguise. Even sentences that naturally favor a generic reading can sometimes elicit an existential interpretation. For example, *I saw birds fly yesterday* most naturally means that I saw *some* birds fly yesterday.

A universal generalization can also masquerade as a generic. For example, someone might ask, *What distinguishes all the green bottles in the bin from the clear ones?* And I might respond, *Green bottles have narrow necks*. The most natural interpretation of my response does not belong in an encyclopedia entry about members of the kind as such; rather, it characterizes all the green bottles in a very small and not particularly special domain. Because this bit of quantitative information is sufficiently clear, I can elide a sentence-initial use of *all*, *the*, or *all the*.

The thought that, on its generic interpretation, *Ks are F* is about Ks as such is more elusive than one would like. But the contrast between (1) and *Tigers are on the front lawn*, and between (9) and *Green bottles have narrow necks*, is highly suggestive. Furthermore, we seem to have an understanding of the sort of generalization that would belong in an encyclopedia entry about Ks—one that we can provisionally work with. Sufficient clarity on the matter requires a theory of the subject we are now trying to demarcate.[6]

(1)-(10) vary in strength. (1) and (2) appear to be roughly equivalent to *Typically, tigers have stripes* and *Generally, birds fly*. (3)-(7) are much weaker, as the following considerations indicate: only a queen bee and her drones (roughly, 15% of the hive) have reproductive capabilities; there may not be a single Orange-Crusher that can crush oranges, or that ever has; perhaps all firefighters are untested novices; less than 1% of mosquitoes carry the West Nile virus; and the overwhelming majority of pit bulls are warm, affectionate, and harmless—nevertheless, there is a defeasible presumption in favor of

[5] See Leslie (2012, p. 354), Nickel (2016, p. 13), and Collins (2018, pp. 35 and 43).

[6] When I say that generics are about Ks as such, I do not mean that they are about Ks *intrinsically*. What I mean is that they are about Ks *unqualifiedly*. Qualification may well be possible—it may even forestall misunderstanding—but truth does not hinge on it. Leslie and her collaborators seem to express the idea by distinguishing "category-wide" generalizations from those with restricted domains (Leslie et al. 2011, p. 17).

thinking that (3)-(7) are all true (Almotahari 2022*a*). Statistically, (8) and (9) are quite strong. Their truth seems to require that almost all members of the kind instantiate the relevant property. But (10) is the strongest of the bunch. Even if every Supreme Court Justice happens to have an even Social Security number, the most natural reading of (10) would be false.

Sterken (2015*c*) argues that, with a bit of contextual prodding, (10) can express a truth. Imagine that two friends are hosting a party at which all the guests are required to have even Social Security numbers. "In the hopes of providing information that will help compile the list of guests, one of the friends says (10). In a context where all Supreme Court justices have even Social Security numbers, my informants hear (10) as true" (p. 7). I am not sure what to make of these judgments. My own reaction is very different: even here it seems that, if (10) is a claim about Supreme Court Justices *as such*, then its use is objectionable. My informal survey of students and colleagues is nowhere near as univocal as Sterken's. Those who judge that (10) is acceptable find some way to hedge or qualify their reaction. When offered an alternative that clearly lacks a generic reading, like *All the Supreme Court Justices have even Social Security numbers*, they favor it over (10). This suggests that the interpretation of (10) on which it is true is akin to the unproblematic context-bound interpretation of *Green bottles have narrow necks*. Both are about a particularly narrow and unremarkable subset of Ks, not Ks as such. So a theory of generics needn't predict the reading.

Consider a slightly different case. Suppose we go to a petting zoo. The proprietor might welcome us by saying, *Feel free to wander around. We've got bears, wolves, giraffes, and tigers. Bears are tame. So are wolves and giraffes. Tigers, though, are fierce. Better leave them alone for now.*[7] In this context, I do not hear the sentence, *Bears are tame*, as a generic about bears as such, but as something basically equivalent to *All the bears are tame*, where the phrase in subject position is contextually restricted to mean all the bears in the petting zoo. Plausibly, the quantifier has been elided from the sentence.

If one thinks (10) is true in Sterken's party-invitation example, it is because the background she provides is enough to make some relevant quantitative information sufficiently clear to justify eliding a sentence-initial occurrence of *all the*. This account has a notable virtue. It explains why the informants who seemed to confirm Sterken's judgment preferred the use of *All the Supreme Court Justices have even Social Security numbers* to (10). The preferred sentence expresses the asserted content more perspicuously. I suspect that other cases are susceptible to similar treatment.

One more example, for good measure:

(11) Prime numbers are odd.

It is common knowledge that 2 is an even prime number, so (11)'s most salient reading is widely judged to be false. But Sterken observes that "there are contexts in which (11) does express intuitively true generalisations. Consider a context where a student is looking at a blackboard with numbers on it, and that student is looking for primes amongst the numbers on the board. A helpful onlooker remarks (11).

[7] Cf., Nickel (2012, p. 296; 2016, pp. 15-16).

In such a context (11) is intuitively true" (2015*c*, p. 22).[8] If we grant the intuition, then it is not clear the example has any significance for a theory of generics, since it is obviously about the subset of prime numbers represented on the blackboard, not prime numbers as such. If we stipulate that the helpful onlooker was making a claim about primes as such, the intuition vanishes.

A theory of generics is not answerable to every interpretive possibility associated with the form *Ks are F*, just as a theory of questions is not answerable to every interpretive possibility associated with *wh*-expressions. Typically, "echo-questions" are excluded from the target domain (Dupre 2021). How, then, should we delimit the range of interpretations to which a theory of generics *is* answerable? This brings us back to our initial challenge: what is a generic sentence?

One often gets the impression that contextualists want a theory of generics to account for every reading that a sentence of the relevant surface form elicits, as if a theory of generics were nothing but a theory of bare plurals.[9] It is not surprising, then, that they embrace eliminativism about genericity—a view according to which there is no such thing as distinctively generic generality. The surface form of (1)-(10) does not pin down a semantically or metaphysically uniform phenomenon. And what little uniformity may initially seem present quickly disappears when we step back and take a broader look. *Tigers are on the front lawn* and *Green bottles have narrow necks* exemplify the relevant form, after all, but they lack the characteristic umph of generic generality.

An alternative method dictates that we be more discerning and take only a proper subset of the possible interpretations associated with the relevant grammatical form to define our subject. A theory of generics stands in a more complicated relation to a theory of bare plural constructions than contextualists admit. If we take the subject of our inquiry to be the linguistic expression of a form of generality that characterizes Ks as such, rather than a highly restricted subset of them, then I think the proper reaction to semantic heterogeneity is not eliminativism about genericity, but greater care in identifying instances of the target phenomenon.

Given this alternative method, we should begin with genericity: it is a form of generality that characterizes the members of a kind as such, and that has a characteristic inductive/explanatory/practical significance, as exemplified by the most salient interpretations of (1)-(10). A generic reading of a sentence is one on which it expresses genericity. A generic sentence, then, is a linguistic construction that naturally elicits a generic reading. These assumptions are not mandatory. We may have to give them up. For now, I suggest we treat them as helpful rules of thumb. Theorizing has to begin somewhere. I choose to begin here. And beginning here complicates the standard case for contextualism; for in addition to showing that a sentence of the form *Ks are F* varies in truth-value from one context to another, the contextualist must show that in each context the relevant sentence expresses a generalization about Ks as such. I submit that many contextualist arguments—particularly those that motivate

[8] I adapted the numbering in this passage to fit my presentation.

[9] This impression is very strongly conveyed (unintentionally, no doubt) in Nguyen (2020) and Lee and Nguyen (2021). In the former, the author says, "I use 'bare plural' to refer to generic sentences of the form *Ks F*..." (p. 1304).

eliminativism—fall short of satisfying this further constraint. I will discuss the most prominent examples as I press on, but there are far too many for one paper to discuss them all.[10]

2. *Omne Trium Perfectum*

Leslie's theory comprises three core claims. The first is primitivism: we possess a fundamental cognitive disposition to generalize information about kinds—a "primitive projective propensity", as Nickel (2016) calls it—the manifestation of which plays a crucial role in the production and evaluation of generic sentences. This means that comprehending a generic generalization involves a default pattern of thought—one whose instantiation is automatic, effortless, and fast. Furthermore, the capacity to engage in this pattern of thought is an innate cognitive endowment. We do not learn how to exercise it by weighing reasons and making inferences; we acquire the knowledge as a result of normal human development. Finally, the interpretation of a generic sentence engages the cognitive mechanisms that underlie our primitive projective propensity. Leslie identifies these mechanisms with the system responsible for "thinking fast", namely, System 1 (Kahneman 2011). Notoriously, the behavior of this system is constrained by various heuristics and biases. Their influence can be overridden with care and concentration. But even the wisest among us are not immune to their influence.

Now that we are clearer about the sort of work a psychologically primitive mode of generalization is supposed to do, we can return to the question I raised in the introduction. Why might one think that primitivism and contextualism are incompatible?

A primitive projective propensity is, in many ways, really stupid. It relies on System 1, after all! Very young children deploy it at roughly the same level of competence as adults. And yet, if it can handle the many examples that allegedly motivate contextualism, it must also be very smart. A good deal of intelligence is needed to reason appropriately about such a diverse and seemingly open-ended range of contextual factors. Positing a cognitive mechanism that transitions between stupid and smart in just those circumstances where it must feels a bit like cheating. Here we have a compelling argument for incompatibility—not one that yields a formal contradiction, but one that endangers the plausibility of conjoining the two doctrines. Regrettably, my response to this argument will not be as compact. But I am able to summarize its general shape before I proceed.

Sentences of the form *Ks are F* are context sensitive, but not every interpretive possibility for a sentence of that form expresses genericity, and only the ones that do are generic sentences properly so-called. This was the central take-away of the previous section. My claim here is that this understanding of our subject vitiates the case for open-ended ("radical") context sensitivity. If the kind of context sensitivity to which generics are susceptible is ubiquitous, unique, but of a highly constrained sort, then the default cognitive mechanism that subserves their interpretation needn't be especially smart to handle it. The paradigm here is the context sensitivity of demonstratives (Kaplan 1989; cf., Nickel 2012). The claim that demonstratives exemplify a ubiquitous and unique sort of context sensitivity is, I presume, compatible with the hypothesis that linguistic interpretation is driven by a primitive psychological

[10] I am not now trying to refute the eliminativist. I am trying to identify the subject of my inquiry and the way someone engaged in that inquiry ought to react to the case for eliminativism.

mechanism. In fact, several prominent accounts of language processing crucially rely on this presumption (Fodor 1983; Borg 2004; Harris 2022). A Kaplanian theory of demonstratives is no threat to a conception of language production and interpretation as modular.[11] The reason why is that demonstrative context sensitivity is highly constrained. So if I can motivate the idea that generic context sensitivity is not as undisciplined as radical contextualists maintain, we can reasonably dismiss the charge of foul play.

Leslie's theory consists of three core claims. The first was primitivism. The second is about logical form. Generics instantiate a familiar tripartite quantificational structure (Heim 1982). Abstracting from irrelevant complexity, the structure is roughly this: Gen $x_1...x_n$ Restrictor $x_1...x_n$ Scope $x_1...x_n$. There is no known language in which *Gen* is associated with an element of surface grammar. Universally, it is an unpronounced syntactic constituent. This idea plays an important role in several subsidiary claims: that *Gen*'s semantic contribution cannot be specified in purely set-theoretic terms; that the relation between *Restrictor* and *Scope* on which the truth of a generic hinges is fixed by our default mode of generalizing; and that the formal representation of *Gen*'s semantic contribution is disquotational. These subsidiary claims elaborate and depend on the more fundamental commitment to *Gen*'s psychological reality.

The final core claim of Leslie's theory specifies the conditions in which a generic of the relevant form is true. This is not a claim about how a competent user of the language computes or represents the meaning of the target sentence; it is a metaphysical theory of the nature of genericity.

> *Ks are F* is true iff the counterinstances are negative, and:
> If F lies along a characteristic dimension for the Ks, then some Ks are F, unless K is an artifact or social kind, in which case F is the function or purpose of the kind K;
> If F is striking, then some Ks are F and the others are disposed to be F;
> Otherwise, almost all Ks are F.

Each disjunct is meant to explain a different type of case. Characteristic-property, artifact-kind, and social-kind generics, like (1)-(5) and (10), fall under the first disjunct; striking-property generics, like (6) and (7), under the second; and high-prevalence generics, like (8) and (9), under the third. One context may elicit an interpretation on which a sentence is a characteristic-property generic, while another context may elicit an interpretation on which the very same sentence is a high-prevalence generic.[12] The

[11] System 1 is not a Fodorian module; it is not informationally encapsulated, neurologically localized, or domain specific. But modules are supposed to be innate, automatic, and fast. The argument above relies on these similarities.

[12] Consider:

(12) French people eat horse meat.

In a context where the subject is the characteristic dietary habits of the French, the sentence is intuitively true. In a context where the subject is unhealthy eating patterns prevalent in France, its negation seems intuitively true: *French people eat croissants and baguettes, not traditional foods like horse meat* (Krifka et al. 1995, pp. 81-83;

disjunctiveness of Leslie's account is, therefore, one potential source of content variation across different contexts of utterance. But I will not lean heavily on this observation going forward. Given the nature of my conciliatory project, I will ignore all but the first clause of Leslie's account. The discussion below will be about the distinction between negative and positive counterinstances. The proper way to understand this distinction is conditioned by a primitivist theory of generic thought. Sentences articulating this form of thought express the default generalizations of System 1, so their truth or falsity turns on the way this cognitive system behaves. I will argue that its behavior permits a great deal of context sensitivity.[13]

3. Symmetry

Although the core claims of Leslie's theory are largely independent of each other, they form a unified whole that is greater than the sum of its parts. The purpose of this section is to convey an understanding of this unity, and to criticize one admirable attempt to coopt some measure of it. If theories that posit open-ended context sensitivity have difficulty accommodating the evidence in support of primitivism, then we have an empirical basis for supposing that generic context sensitivity is highly constrained.

The first core claim of Leslie's theory (that generics articulate default generalizations) partially explains the second (that *Gen* is a covert feature of logical form). If an engineer were designing a signaling device to encode and send information to a receiver that she knew was normally disposed to interpret certain signals in certain ways, and if the engineer wanted to maximize the signaling efficiency of her device, then she would design it in such a way as to mark the occasions in which the receiver should *diverge* from its default mode of interpretation. Producing discernible marks to signal that the default interpretation is appropriate would, under the circumstances, be unnecessary and therefore inefficient, since the default interpretation will normally be assigned anyway, and the production of a discernible mark would require some time and energy. Assuming that the externalization of language is, to some degree, an efficient way of communicating information, the absence of an overt generic operator is

---

Streken 2015*b*, p. 2504 and fn. 11; 2015*c*, p. 7). But this shifty truth-value assessment can be explained in primitivist-friendly terms. In the first context, (12) is a characteristic-property generic, so the corresponding clause in Leslie's truth-conditions is the relevant one. This means that, if there are no positive counterinstances, then the sentence is true if horse meat lies along the relevant dimension and at least some French people eat it. Obviously, the French characteristically consume many kinds of food, but consuming one kind does not preclude the consumption of other kinds. So the foods are not incompatible alternatives and thus not positive counterinstances. Since horse meat is part of a characteristically French diet, (12)'s truth in the first context is unproblematic. In the second context, (12) is a high-prevalence generic. This means that a different clause in Leslie's account comes into play, specifically, the clause that requires a high prevalence of French people to eat horse meat. But, given the way the case is set up, we are supposed to assume that relatively few French people eat horse. So, relative to this context, (12) is a false high-prevalence generic. Its negation is therefore true.

[13] I learned of Plunkett et al. (2023) only after this paper had been submitted for review. It presents a novel critique of Leslie's theory but will not be discussed here.

explained by the fact that generic sentences engage a default mode of interpretation for which overt marking is unnecessary. It is not surprising, then, that no known language articulates *Gen*.[14,15,16]

Furthermore, if our most basic system of thought about kinds facilitates the production and interpretation of generic sentences (the first claim), then it is not surprising that, across a wide range of cases, generic sentences are assigned truth-values in a somewhat messy way (the third claim). From an adaptationist standpoint, basic systems of thought are shaped by local selection pressures and largely haphazard genetic constraints. These systems persist because they worked well enough in the conditions where they were put to use by our ancestors. There is no *a priori* reason to believe that, under these conditions, our basic systems should evolve sensitivities that are systematic, non-disjunctive, or highly "joint-carving". In fact, if it were otherwise, we would expect there to be an explanation. Evolutionary debunking arguments appear to be based on a similar sort of expectation: what explains the fact that our beliefs about morality, mathematics, and ordinary medium-sized objects are largely true or joint-carving, given that the cognitive systems that produce them were blindly shaped by evolution? This familiar challenge presumes that evolved systems behave in somewhat quirky ways. It is deviation from the norm, after all, that cries out for explanation. Anyone familiar with the literature on generics knows that our truth-value judgments about them appear quirky. Different theories handle this in different ways. One nice feature of Leslie's theory is that the quirkiness is a fairly straightforward consequence of its primary core thesis and standard evolutionary reasoning.

Finally, the second and third claims jointly support the first. Children acquire an understanding of generics at roughly the age of two (Gelman 2003). But how is that possible, given (i) the truth-conditional complexity of such sentences (the third claim) and (ii) the absence of *Gen* from their surface form (the second claim)? This question is made even more challenging by two additional observations: (iii) children find it immensely difficult to associate information with absences; and (iv) although quantifiers are far more theoretically tractable than *Gen*, "recent developmental findings suggest that

---

[14] The assumption that linguistic communication is efficient plays a crucial role in an ongoing interdisciplinary research program. Summarizing many decades of influential work on a wide range of subjects (the semantics of color terms and logical connectives, ambiguity, syntactic dependency relations, and compositionality) Gibson et al. conclude that "across levels of linguistic analysis, from words to syntax, the form of human language exhibits a strong tendency to be structured for efficient use" (2019, p. 389). This is to say that the assumption is by no means *ad hoc*.

[15] *Objection*: If generic generality is the default, as the primitivist maintains, why should there be a syntactic element dedicated to it at all? (Thanks to Gabe Dupre for raising this worry.)
Reply: Communicative efficiency constrains only the features of a sentence that play an essential role in externalization, like its phonological profile or surface structure, not the features that belong to it by virtue of Universal Grammar. As I understand it, Leslie's view is that *Gen* is a feature of generics by virtue of UG.

[16] Another option might be worth exploring. According to the Minimalist Program, language design is computationally optimal (Chomsky 1995). This idea is taken to imply a few different "principles of *minimal computation*", one of which governs the way sentences are phonologically realized: "pronounce as little as possible" (Chomsky 2017, p. 31). Given this general constraint on the design of language, and the default status of generic generalization, the fact that *Gen* is universally unpronounced might be explained by computational optimality rather than communicative efficiency.

generics may be as easy as quantifiers for young children to acquire and process, and in some cases even easier" (Leslie et al. 2011, p. 16).[17] These observations generate "the paradox of generic acquisition" (Leslie 2008). Its solution, Leslie maintains, is that generic thought is psychologically fundamental: it is an innate, default mode of generalization the outputs of which are clothed by our language faculty in the form *Ks are F*.

In light of all this, Leslie's theory exhibits an elegant symmetry.

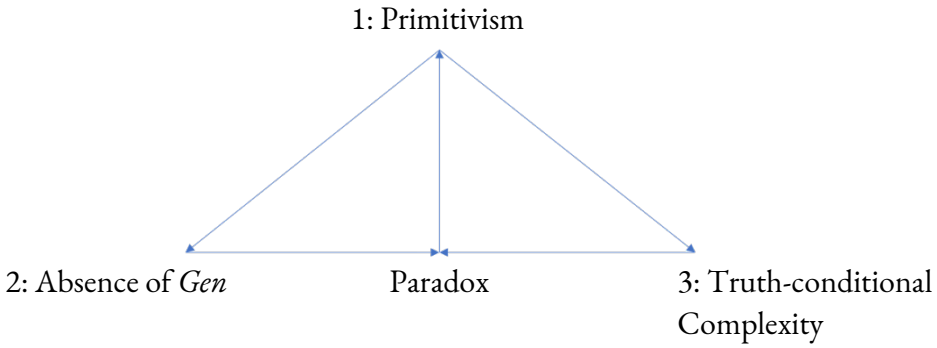

FIGURE 1: The structure of Leslie's theory.

The mutually reinforcing theoretical relations between the core claims (represented in the diagram by arrows symbolizing directions of support or explanation) is what I was alluding to when I said that Leslie's theory enjoys a unity that makes the whole greater than the sum of its parts.

Admirably, Sterken provides a competing account of both the absence of *Gen* from surface grammar and the paradox of generic acquisition. In each case, her explanation is given in terms of the hypothesis that *Gen* is a covert indexical.

> On the indexical approach, it is somewhat unsurprising that *Gen* is unpronounced: *Gen* is an instance of a certain type of expression in natural language which is frequently unpronounced—*Gen* is a covert indexical which is represented as a free variable at the level of logical form, and several such expressions are never pronounced. ...Concrete examples include quantifier domain variables and implicit argument places (2015*a*, p. 24).

The datum requiring explanation is that *Gen* is a covert syntactic constituent. I fail to understand how the hypothesis that *Gen* is a covert indexical explains that datum. There is not enough distance between the datum and the hypothesis for the relationship between them to qualify as explanation. Suppose *Gen*

[17] As Leslie and her collaborators note, "The results of the main experiment were thus not due to a basic lack of competence with the quantifiers; they seemed instead to do with the difficulty of processing category-wide quantified statements. When confronted with a quantified claim about an entire category (as opposed to a specific subset), young preschoolers appear to rely on their interpretation of the corresponding generic" (p. 17). Crucially, it is not Leslie's view that, prior to the age of 4, children do not understand universal or existential quantification.

is the type of expression whose instances are, as Sterken says, "frequently unpronounced"; still, some tokens of this type frequently are pronounced. Why, then, is *Gen* not at least sometimes an overt syntactic constituent? Responding that it's a *covert* instance of the type relies on the very thing in need of explanation.

Perhaps my reaction is unfair; maybe the idea of indexicality does shed light on why *Gen* is covert. Some of Sterken's remarks suggest as much: "If *Gen* is construed as an indexical (or free variable), it is at least unsurprising that it is never pronounced" (p. 24). My understanding of indexicality is exhausted by the idea of rule-governed variation in content across different contexts. This idea is fundamentally semantic. But the datum to be explained is a truth of syntax. *Prima facie*, this difference suggests an explanatory gap (Chomsky 1975). Perhaps some syntactic phenomena have semantic explanations—NPI-licensing is the first thing that comes to mind (Hintikka 1977)—but in this case I fail to see why indexicality would make covertness less surprising.[18] Is Sterken suggesting that *Indexicals are covert* is a true generic? If it is false, then I do not see how we have an explanation here. But I do not think we have any reason to believe that it is true. There is simply nothing exceptional about overt indexicality.

It seems that Sterken is assimilating the question, *Why is* Gen *covert?*, to the question, *Why are some free variables covert?* If indexicalism is true, then the former is just a special case of the latter, to which many of us owe an answer anyway. Consolidating one's explanatory debts is progress. But it is not an explanation. And there is reason to be worried about indefinite postponement; for, as Sterken is aware, "...no one has a good story about why [domain variables, implicit argument places, *PRO*, and existential closure] are not pronounced" (2016, p. 525).

Does the hypothesis that *Gen* is a covert indexical provide an equally promising account of the acquisition paradox? Here, again, is Sterken:

> Children do not need to learn what *Gen* means since it has no fixed meaning. *Gen* is like other supplementives—for example, quantifier domain variables, implicit arguments places (e.g., the standard or comparison class variables of the predicates *tall* and *smart*), and demonstratives. Minimally, then, the requisite abilities to acquire generics are the cognitive or conceptual ability to generalise in some way, and the ability to saturate or resolve the value of free variables (of the appropriate semantic type) in some way. ...Whereas Leslie's solution to the Paradox is grounded in the cognitive mechanism of primitive, default generalisation, the solution of the indexical approach is plausibly grounded in the linguistic properties of generics and the mechanism of saturating an indexical (supplementive) (p. 26).

Two points are relevant. First, according to the indexicalist, competently saturating *Gen* is a matter of assigning the appropriate quantificational force (universality, typicality, existence, etc.). Presumably, this requires the ability to competently process category-wide quantified statements, not just quantified statements about a small and unremarkable set of items (e.g., all the crayons in the box over there), since

[18] For a persuasive response to Hintikka, see Chomsky (1980*b*, pp. 123-125).

Ks are F is about members of the category as such.[19] But it is precisely this sort of competence that two-year-olds lack, even though they are able to understand category-wide generics (Leslie et al. 2011). So the indexicalist's response to the acquisition paradox seems empirically untenable. Second, there is a datum closely related to the paradox that the indexicalist seems unable to explain: the acquisition of competence with category-wide quantificational sentences basically coincides with the onset of inhibitory behavior at roughly the age of 4. Why should that be? Leslie's account provides an explanation: understanding the meaning of these sentences requires inhibiting the activation of one's default mode of generalization; the capacity to inhibit default action is not acquired until roughly age 4; so we should expect kids to appreciate the meanings of category-wide universal and existential sentences at about the time they can inhibit default behavior. From the indexicalist's point of view, this pattern of development is merely coincidental.

Leslie's theory has been extended in various directions. For example, it sheds light on why a certain fallacy (namely, "the generic overgeneralization effect"—our tendency to conflate a strong statistical generalization with a corresponding generic) is both pervasive and difficult to resist (Leslie et al. 2011); it can account for one possible source of prejudicial thought (Leslie 2017); and, combined with a promising take on noun-phrase ambiguity, the theory can explain the normativity of certain generic sentences, like *Boys don't cry* and *Women are nurturing* (Leslie 2015*b*). These extensions have been challenged, but the theory is a crucial part of a thriving interdisciplinary research program.

Two of the applications I mentioned in the previous paragraph will be relevant for the discussion below. Drawing attention to them here will make things somewhat easier a little later.

Generic overgeneralization occurs when a true statistical claim is understood as a false generic. The generalization might, therefore, seem more credible in light of how much support the statistical claim enjoys (Cimpian et al. 2010; Leslie et al. 2011; Sorensen 2012; Almotahari 2022*b*). I recently had a conversation with a well-educated native speaker of English who very briefly maintained that *Bees are sterile* is true! When this person was reminded of queen bees, they sheepishly smiled. All it took for this person to assent to the false generic was the salience of the statistical truth that the vast majority of bees are sterile.

Leslie's account of normative generics relies on an independently motivated theory of nouns. For example, *woman* expresses a "dual-character" concept, the applicability of which depends on two distinct (though related) sets of criteria—one descriptive, the other normative (Knobe et al. 2013). Something satisfies the normative criteria only if it lives up to the social ideals associated with the kind. This explains why someone who says of Hillary Clinton, *She isn't a real woman*, does not evince *semantic* incompetence. The speaker is merely signaling that the relevant criteria for membership in the kind *woman* requires the exemplification of certain "virtues", e.g., vulnerability, deference to authoritative men, and the absence of political ambition. Normative generics range over the set of Ks that exemplify the relevant set of ideals. When conversational partners endorse the ideals, treating them as principles that demand conformity or assent, *Ks are F* acquires a normative reading.

[19] Denying this presumption would make the competent saturation of *Gen* considerably more mysterious. It would mean that although two-year-olds are able to saturate *Gen*, they are not able to understand the corresponding category-wide generalization. I am not sure this is even coherent.

4. Semantics, Leslie Style

The central claim of this section is that Leslie's truth-conditions predict a good deal of context sensitivity unique to generics. A subsidiary claim is that many of the alleged counterexamples to her metaphysical semantics can be dealt with. I am unable to solve every problem in the vicinity, only enough of them to warrant the assumption that Leslie's truth-conditions provide an illuminating (if not exceptionless) *model* of genericity—one that I will use to vindicate the core thesis of this paper: that primitivism and contextualism are compatible.

It is one thing to show that a picture of genericity makes false predictions; it is quite another to show that it is a bad model. The latter depends in part on our theoretical aims. My aim is broadly logical: I want to demonstrate the compatibility of two doctrines. For this purpose, Leslie's truth-conditions seem particularly well suited, since they are framed in a way that presupposes a primitivist take on the processing of generic sentences. If I can show that her truth-conditions entail an empirically viable and unique form of context sensitivity, "compatibilism" will have been vindicated.

A negative counterinstance to a generic is a member of the relevant kind that lacks the relevant feature, but not by having an incompatible alternative. The clearest examples involve absences. Stripeless tigers are negative counterinstances to (1); flightless birds are negative counterinstances to (2); and sterile worker bees are negative counterinstances to (3). In each case, the feature that the corresponding generalization predicates is simply uninstantiated. But negative counterinstances do not always involve a mere absence. In many cases, they involve the presence of an incompatible alternative—one that is not particularly vivid, memorable, or salient. It is *as if* the alternative feature were absent. So, a positive counterinstance to *Ks are F* is a member of the kind K that exhibits an incompatible alternative to F which is at least as vivid, memorable, and salient as F is (Leslie 2008, pp. 33-36). The difference between negative and positive counterinstances "is not intended as a metaphysical distinction, but rather a psychological one. What matters is whether we *take* the counterinstances as negative or positive. …human perception is the important factor" (p. 34, my emphasis).[20]

Here is how Leslie applies the idea:

> Generics such as 'peacocks have blue tails' and 'lions have manes' are readily judged true because the females of the species lack these properties; they do not exhibit competing properties, *or at least the competing properties they do exhibit are quite boring and unremarkable*. …There is an intuitive difference between simply lacking a

[20] Sterken maintains that flightless birds are positive counterinstances to (2), since they rely on alternative modes of locomotion: walking, swimming, and running (2015*b*, p. 2497). But *walks*, *swims*, *runs*, and *flies* are ambiguous. On one interpretation, they describe a specific episode of movement (e.g., *Yesterday, I saw Xavier walk*.) On another interpretation, they describe stable habits of movement. (*What are some of Xavier's hobbies*? – *Well, he walks, runs, and swims*.) On the first ("stage-level") interpretation, it's plausible that nothing can walk, run, swim, and fly. But on the second ("individual-level") interpretation, a normal mature duck can walk, swim, run, and fly. So the corresponding properties are not incompatible alternatives and thus not positive counterinstances. (2) expresses the latter reading.

> feature and lacking it in virtue of having another, *equally memorable*, feature instead (p. 35, italics added).

It is partly because the boring brown tails of female peacocks are not as vivid, memorable, or salient as the fabulous blue tails of male peacocks that *Peacocks have blue tails* is true. One aspect of this quotation requires emphasis: the factors that make for a positive counterinstance must enjoy some degree of psychological stability. Salience and vividness are not enough if they are fleeting. That explains why Leslie unpacks the idea in terms of being "equally memorable". In the absence of this stability, an alternative trait will not constitute a positive counterinstance.

Vividness, salience, and memorability are properties that a feature might have in one context but not in another. Something might be particularly vivid, salient, and memorable when we discuss it, but not when others do. If we happen to be especially knowledgeable about the subject of peacock tails, or highly pedantic, or just curious about female peacock anatomy, then *Peacocks have blue tails* is unlikely to express a truth relative to our context. The otherwise unremarkable and easily forgotten brown tails would probably qualify as positive counterinstances.

Here is a more concrete way of implementing the idea—one that I will eventually trace back to Leslie's original discussion: in a typical context where the sentence, *Peacocks have blue tails*, expresses a truth, the relevant psychological factors determine a contrast class including salient alternatives like having purple tails, or having green tails, etc. Crucially, the property of having a brown tail will not be a member of the relevant contrast class. As a result, *Peacocks have blue tails* will express the generalization that peacocks have blue tails rather than purple tails or green tails—and that is true. In a context where the relevant psychological factors determine a contrast class including the property of having a brown tail, the sentence will express the generalization that peacocks have blue tails rather than brown tails—and that is false.

The point is perfectly general. Leslie's view predicts that *Ks are F* is elliptical for the generalization that Ks are F rather than G. To fully specify the generalization that, in context, serves as the content of the generic, one must identify the contrast class that salience, vividness, and memory make available to the speaker and her audience. By accommodating the "inherently contrastive" aspect of a generic (Leslie's expression) and tying it to the highly context sensitive presence of psychological factors like salience, vividness, and memory, Leslie's theory predicts that generics are always context sensitive. Remember, on the view in question, the truth of a generic *always* hinges on whether the counterinstances are negative, and (as we just saw) whether a counterinstance is negative *always* hinges on contextually variable psychological factors.

The factors that explain the negative-positive distinction are one source of context sensitivity. These psychological factors play no role in the semantics of statistical quantifiers. This is important because it means that, given Leslie's view, we should expect the context sensitivity associated with generics to be distinctive! I emphasize the point because Sterken argues at great length, and with considerable ingenuity, that distinctive context sensitivity tells in favor of treating *Gen* as an indexical

(2015*b*, p. 2504).[21] However, if Leslie's account predicts distinctive context sensitivity, the case for indexicality is jeopardized. As I will argue momentarily, Sterken's examples can receive an adequate explanation in terms of the negative-positive distinction, or the disjunctiveness of Leslie's truth-conditions, or the mechanisms through which contextual effects are generally explained.

The context sensitivity that Leslie's account predicts is often neglected. This is an understandable mistake. The official statement of Leslie's metaphysical semantics does not include an explicit reference to contrast classes. But Leslie does draw attention to the idea in various places (2008, pp. 35 and 40). Her remarks are brief, but they are punctuated with a bit of fanfare: "This contrastive feature of judgments of generics has not been noted in the literature so far" (p. 35). I want to suggest that it undermines an influential objection to primitivism.

> One driving force of Leslie's theory is that generics express our most basic, default generalizations about a kind, that we are hesitant to give up in light of new and conflicting evidence. Thus, in a substantive sense, *the primitive generalisations expressed by generics on Leslie's account are quite coarse and stable over time, and across different contexts* (Sterken 2015*b*, p. 2507, emphasis added).

One claim that Sterken appears to be making here is that if the meaning of a generic sentence were a default generalization, then generics would not be context sensitive. This claim is reiterated in Nguyen (2020, p. 1308) and in Lee and Nguyen (2021, p. 1757, fn. 25). If the claim were true, then evidence of context sensitivity would be evidence against primitivism. But the claim is not true.

The context sensitivity that I have been spotlighting undermines some other important criticisms. For example, Sterken asks us to consider the following sentences:

(13) Mammals lay eggs.
(14) Mammals give birth to live young.

[21] Sterken was most concerned with arguing that a disquotational theory of *Gen* is incapable of accommodating the sort of context sensitivity her examples were supposed to demonstrate. In contrast, I am developing a view on which the source of context sensitivity is our default mode of generalization's receptivity to different salience relations. This would take some of the pressure off the semantics of *Gen*. Why think this aspect of System 1 cognition is variable in the way I am supposing? Why is the relevant salience relation not fixed and innate? I am grateful to an anonymous referee for prompting me to address these questions more directly. Part of the answer has already been given. Salience is probably an important factor in the resolution of demonstrative reference and yet there is compelling reason to believe that the language faculty is modular (innate, automatic, and so on). In fact, one important respect in which modules differ from System 1 is noteworthy: whereas modules are informationally encapsulated, in that they are controlled by a relatively small body of information that is closed off from one's overall knowledge about the world, System 1 is informationally porous. Since salience is determined at least in part by what one knows, as new information permeates System 1 over time, it would be reasonable to suppose that salience relations are not permanently fixed. See Section 5 for further discussion.

She acknowledges that Leslie's truth-conditions correctly predict the falsity of (13). The mammals that give birth to live young are positive counterinstances. However, Sterken goes on to claim that Leslie's truth-conditions incorrectly predict the falsity of (14), since platypuses and echidnas are positive counterinstances (2015*b*, p. 2497). But egg-laying mammals (aka "monotremes") are not as vivid, memorable, or salient (outside Oceania) as the mammals that produce live young. Many of us do not perceive platypuses and echidnas as positive counterinstances to (14), and "human perception is the important factor" (Leslie 2008, p. 34). Even if monotreme birthing is made salient, its salience does not seem to enjoy the psychological stability that is required of a positive counterinstance.[22] That is why, for many readers, the intuition that (14) is true persists in our context. None of this endangers Leslie's account of (13), however, because mammals that give birth to live young plausibly enjoy the stable sort of salience that is required for being perceived as a positive counterinstance. Interestingly, the small group of specialists I talked to, for whom platypus and echidna biology are highly salient, find both (13) and (14) objectionable, preferring statistical generalizations or a more complex construction: *Mammals give birth to live young or lay eggs*. But this nicely coheres with Leslie's truth-conditions. More knowledge of taxonomic complexity is likely to correlate with differences in psychologically stable salience relations (van Rooij MS).

Although Sterken acknowledges that (13) is false, she claims that its occurrence in (15) can, in the right environment, express a truth. Just imagine a context where a zoologist is lecturing about birds and their relationship to other species (2015*c*, p. 8).

(15) Birds lay eggs. Mammals lay eggs, too.

How might Leslie explain this?

Before I sketch an explanation on Leslie's behalf, consider a slight variant of Sterken's example:

(16) Characteristically, birds lay eggs. Some mammals lay eggs, too.

The follow-up sentence (*Some mammals…*) is weaker than the initial sentence (*Characteristically…*), which is basically synonymous with the generic reading of *Birds lay eggs*. This demonstrates that the felicity of (15) does not require *Mammals lay eggs* to be a generic sentence. Now, we saw that instances of the form *Ks are F* can elicit existential readings. Recall *Tigers are on the front lawn* and consider *Yesterday, I saw mammals lay eggs*. In fact, discomfort with a generic sentence often triggers interpretive repair that settles on a nearby existential generalization. In my experience, this form of charitable reinterpretation is quite common; it plausibly explains why (13) is most naturally understood as an existential in *Yesterday, I saw mammals lay eggs*. One possible explanation of Sterken's example, then, is this: we imagine an expert uttering (15); we naturally assign an interpretation to the utterance that

[22] While lecturing to a class of roughly 40 third-year philosophy majors, I informed the students that platypuses are mammals and that they reproduce by laying eggs. I then asked whether it is true that mammals lay eggs. Some expressed uncertainty and many said no. Not a single student said yes. This anecdote is not dispositive, but it is not negligible either. It is almost as if the students did not believe that platypuses are mammals.

does not attribute an amateurish mistake to the speaker; if *Mammals lay eggs* expressed a generic generalization, then by our lights the expert would be making an amateurish mistake—we are assuming, after all, that (13) is false; so we reach for the existential reading. (15) is unproblematic, then, because it is roughly equivalent to (16).

A nice feature of the account I just sketched is that it correctly predicts the felicity of the following sentence:

(17) Birds lay eggs. Mammals lay eggs, too, though they don't do so characteristically.

For the account implies that (17) is equivalent to

(18) Characteristically, birds lay eggs. Some mammals lay eggs, too, though they don't do so characteristically.

However, if the occurrence of *Mammals lay eggs* in (15) were a generic, then (17) would be inconsistent, because on its generic reading *Mammals lay eggs* is about as strong as *Characteristically, mammals lay eggs*, in which case (17) would be equivalent to

(19) Characteristically, birds lay eggs. #Characteristically, mammals lay eggs, too, though they don't do so characteristically.

And that is plainly incoherent. The upshot, I maintain, is that my explanation of (15) on Leslie's behalf is independently motivated. At no point does it resort to special pleading.

(13)-(15) are part of a much more elaborate attack on the idea that generics are sensitive to the distinction between negative and positive counterinstances. Three more examples deserve consideration.

(20) Danes are tall.
(21) Swedes have blonde hair.
(22) Dobermans have floppy ears.

According to Sterken, each sentence is intuitively true despite common knowledge of the positive counterinstances: short Dutch people; brown-haired Swedes; and pointy-eared Dobermans (Sterken 2015*b*, p. 2497).

Leslie's view explains the temptation to endorse (20) and (21). It is widely believed that a large majority of Danes are tall and Swedes blonde. Because of the tendency to systematically conflate statistical claims with generic generalizations, these widely accepted statistical claims are erroneously voiced with (20) and (21). In other words, Sterken overlooks the possibility of generic overgeneralization as the basis for the intuition that (20) and (21) are true. I am inclined to think that this possibility offers a more satisfying account of the intuition behind these sentences, since they bear a striking resemblance to morally objectionable stereotypes: *Pacific Islanders are fat* and *Kurds have big noses*, not to mention

the many familiar generalizations about Blacks, Jews, and other traditionally marginalized groups. I suggest that it is no great cost if a theory forces us to give up intuitive stereotypes about the physical appearance of a group of people. (Surely *sometimes* the point of philosophy is to change the world.) As for generics that are meant to empower a marginalized group—their acceptability needn't always depend on whether they are true. Slogans can serve a sociopolitical or motivational role that warrants affirmation even if the jury is still out, or the circumstances are arranged in such an unjust way as to constitute their falsity. In these cases, they express lofty aspirations rather than promising theories. And this helps explain why stereotypes about the physical appearance of marginalized groups seem so much worse than (20) and (21), even though they are all false: the former traditionally express pernicious aspirations; they are associated with a worldview that promotes violence and subjugation; the latter are not—at least, not as straightforwardly.[23]

(22) is a far more difficult case. The threat it poses to Leslie's account deserves greater emphasis. Originally due to Nickel (2009; 2012), the example motivates the idea that *Ks are F* can be context sensitive even when it is about Ks as such.

Discourse 1: Some dog breeds have evolved to focus on their hearing. These breeds have pointy ears. Dobermans, however, mostly rely on their sense of smell, which is why Dobermans have floppy ears, not pointy ears.

Discourse 2: Welcome to this year's meeting of the Westminster Kennel Club. Once again,

[23] A referee for this journal suggests that it is important for various sociopolitical aims that acceptability hinge on truth. The prevalence and psychological momentum of bigotry can make the facts about race and gender one of the few tools at our disposal for changing hearts and minds. So, they maintain, it would be a weightier cost than I acknowledge if empowering generics about the physical appearance of some people were false. I agree with the suggestion about the importance of truth, evidence, and rational engagement in circumstances where there are few other means available, but this does not indicate a significant cost for my position. If one wants to rationally engage the bigot with facts and argument, then I think it would be a bad strategy to do so in a way that triggers System 1 cognition. "Thinking fast" makes one *more* susceptible to bias and irrational influence, not less. If primitivism is true, and generics engage System 1, then the referee's suggestion indicates just how little would be lost by adopting my view. Rather than relying on claims that potentially facilitate the effects of bias and bigotry, we should appeal to considerations that stimulate System 2, and that requires more precision than generics can provide. Furthermore, an objectionable stereotype can be an effective sociopolitical tool, as well—one that strongly resembles appropriation in the case of slurs. For example, in a powerful speech delivered in 1967, Kwame Ture (aka Stokely Carmichael) said, "Our noses are broad, our lips are thick, our hair is nappy—we are black and beautiful!" Ture was subverting the disposition to infer *Black people aren't beautiful* from *Black people have thick lips*, but his strategy does not concede the truth of the stereotype (at any rate, it does not have to). Rather, it resembles a frustrated speaker who says, without really meaning it, "Okay, so I'm stupid; then why did you ask for my advice?" The initial move—"Okay, so I'm stupid"—is an expression of the other side's point of view (Holton 1997; Recanati 2010; Stokke 2013). More could be said about these issues. Consider, e.g., *Bikers wear leather* and *Farmers have rough hands*. The former attributes a functional property and the latter is false, though perhaps there is a statistical truth in the vicinity that we might imperspicuously express with it.

> we've got a great range of dog appearances. While Labradors and golden retrievers have floppy ears, Dobermans do not. Dobermans have pointy ears.

Dobermans as such are said to have floppy (not pointy) ears in Discourse 1. Dobermans as such are said to have pointy (not floppy) ears in Discourse 2. Consequently, the challenge that (22) presents cannot be resolved in the way that I explained (15). How, then, can it be resolved?

The answer requires two separate lines of thought. First consider Discourse 2. Plausibly, *Doberman* expresses a dual-character concept; it is associated with two separate application conditions, one of which reflects the Westminster Kennel Club's largely conventional idealization of the breed.[24] A specimen exemplifies this idealization only if it has pointy ears. This is why a snooty dog-show enthusiast might say, *That floppy-eared Doberman isn't a real Doberman*, without evincing semantic incompetence. This does not mean that the remark is entirely free from error. Perhaps it manifests a perverse or questionable aesthetic. The issue is neither here nor there. What matters is that this accessible reading of *Doberman* excludes floppy-eared specimens from its extension. As a result, they are not positive counterinstances. So the truth of (22) as a generic sentence is fully consistent with Leslie's analysis. Of course, the mere fact that this story relies on the dual character of the Doberman concept does not entail that (22) is a normative generic. Whether it is normative depends on whether we endorse the Westminster Kennel Club's ideal. But, again, that is neither here nor there.

A somewhat different strategy seems necessary in the case of Discourse 1, since the hypothetical speaker is focusing on the natural state of Dobermans, not some conventional idealization. The strategy I want to pursue begins with the observation that *have* is said in many ways. For example, some Japanese maples *naturally* have green leaves; others might *artificially* have white ones. If I come across a red Japanese maple, I might wonder, *How does it have such a color?* And the answer might be, *It has the color naturally*.[25] Given that Discourse 1 focuses on evolved canine traits—presumably, in response to a question under discussion about the natural features of the kind—the salient interpretation of *have* renders it contextually synonymous with *naturally have*. So, in Discourse 1, *Dobermans don't have pointy ears* means that they do not naturally have pointy ears—and that is true. Artificially pointy-eared Dobermans are not positive counterinstances to *Dobermans naturally have floppy ears*, since they too naturally have floppy ears. Something can naturally have one property and artificially have an alternative. Likewise, pointy-eared Dobermans naturally have floppy ears but artificially have pointy ones. Although pointy-eared-ness is an alternative to floppy-eared-ness, the property of *artificially* having pointy ears is not an alternative to the property of *naturally* having floppy ears.[26]

---

[24] In this connection, see Leslie (2015*b*, p. 121) for some related discussion about the dual-character concept associated with *dog*.

[25] Cf., Hansen (2011, p. 219).

[26] The content one asserts by uttering (22) might be question sensitive (Roberts 1996). A theory of this sort has to contend with Sterken's challenge to pragmatic explanations (2015*c*, p. 11). This challenge relies on the "A-Quantifier Test", which asks whether the insertion of meaning-preserving adverbial quantifiers would produce the same pragmatic effects across a change in context resembling the difference between Discourses 1 and 2. If not, then the source of context sensitivity must be *Gen* itself, as the indexicalist maintains, not a general conversational mechanism, as I suggest. Sterken claims that the test supports indexicalism. But the appropriate A-

My account of (22) preserves its generic flavor while exploiting some independently motivated sources of context sensitivity. This strategy seems viable quite generally. More importantly, it enhances the degree of confirmation that primitivism enjoys. "A theory is directly confirmed by the observations it predicts/explains, and indirectly confirmed by the discharging of its [explanatory] debts" (Dupre 2021, p. 199).

5. Caveats and Conclusions

My project has been conciliatory and constructive. I tried to say enough about the virtues of primitivist contrastivism to encourage further investigation. It would be a problem for my view if the kinds of salience relations it needs were of a sort that System 1 is too stupid to detect. Nothing I said here settles the matter in the way that I require. But there are reasons for optimism.

First, it is important that we not exaggerate the stupidity of System 1. The examples that demonstrate its erroneous behavior also reveal its susceptibility to correction: "A bat and ball cost $1.10 in total. The bat costs $1 more than the ball. How much does the ball cost?" Most of the undergrads Kahneman surveyed at Harvard and MIT say $.10 (Leslie 2007). But $1 - $.10 = $.90, so the correct answer is $.05. Now, is it just me or does repeated exposure to this example change one's intuition? Somehow, I doubt that I'm exceptional. This sort of maturation is grounded in the informational permeability of System 1. Changing behavior in the light of new information is a mark of intelligence.

Second, the divide-and-conquer strategy I pursued does not treat salience as a cure-all. Many of the most worrisome cases for the primitivist are dealt with in other ways: either by re-classification, or by appeal to disjunctiveness, or by relying on generic overgeneralization and noun-phrase ambiguity. The appeals to salience are limited in a way that does not require System 1 to be a magical black box.

Finally, the hypothesis that *Gen* is an indexical has trouble explaining the acquisition data and the absence of *Gen* from surface form. This should make us reluctant to abandon primitivism, and primitivism motivates a type of context sensitivity grounded in generic cognition.

---

quantifier depends on the type of generic claim (22) is being used to make. Recall footnote 12. A sentence might be a generic of a certain type in one context (thus rendering one kind of A-quantifier close enough in meaning for insertion) and a generic of a different type in another context (thereby rendering a different kind of A-quantifier closest in meaning). It seems that, in Discourse 1, sentence (22) is a characteristic-property generic, akin to *Dogs have four legs*. Roughly, it means that floppy-eared-ness is both the specific value of a dimension along which Dobermans characteristically resemble each other and a property that *some* members of the kind instantiate. The A-quantifier that comes closest in meaning is *sometimes*. Now evaluate the corresponding discourse: *Some dog breeds have evolved to focus on their hearing. These breeds have pointy ears. Dobermans, however, mostly rely on their sense of smell, which is why they sometimes have floppy ears.* I submit that it is true. But, in the context of Discourse 2, my account says that *Doberman* expresses a social kind, so (22) is a social-kind generic. Its truth rests on what happens to be the case *conventionally*, by virtue of a socially imposed function. Furthermore, I hear *Dobermans conventionally don't have floppy ears* as true in the relevant discourse: *While Labradors and golden retrievers have floppy ears, conventionally Dobermans don't. Traditionally, they have pointy ears.* The A-quantified sentences that come closest in meaning to the generics at issue differ in truth-value across the two contexts.

Acknowledgments

Thanks to Gabe Dupre, Sarah-Jane Leslie, Anthony Nguyen, Bernhard Nickel, Brian Rabern, Rachel Sterken, Nick Treanor, and two anonymous referees for their helpful comments and questions.

References

Almotahari, Mahrad. 2022*a*. 'Weak Generics' *Analysis* 82: 405-409.

Almotahari, Mahrad. 2022*b*. '*Kalām* and Cognition'. In M. S. Zarepour, ed., *Islamic Philosophy of Religion: Analytic Perspectives*. Routledge.

Borg, Emma. 2004. *Minimal Semantics*. Oxford University Press.

Chomsky, Noam. 1975. *Questions on Form and Interpretation*. De Gruyter Mouton.

Chomsky, Noam. 1980*a*. 'Rules and Representations' *Behavioral and Brain Sciences* 3: 1-61.

Chomsky, Noam. 1980*b*. *Rules and Representations*. Columbia University Press.

Chomsky, Noam. 1995. *The Minimalist Program*. MIT Press.

Chomsky, Noam. 2017. 'Two Notions of Modularity', in R.G. de Almeida and L.R. Gleitman, eds., *On Concepts, Modules, and Language: Cognitive Science at Its Core*. Oxford University Press.

Cimpian, Andrei et al. 2010. 'Generic Statements Require Little Evidence for Acceptance but Have Powerful Implications' *Cognitive Science* 34: 1452-1482.

Collins, John. 2018. 'Genericity *sans Gen*' *Mind* & *Language* 33: 34-64.

Dupre, Gabe. 2021. 'Linguistics and the Explanatory Economy' *Synthese* 199 (Supplement 1): 177-219.

Fodor, Jerry. 1983. *The Modularity of Mind*. MIT Press.

Gelman, Susan. 2003. *The Essential Child: Origins of Essentialism in Everyday Thought*. Oxford University Press.

Gibson, Edward et al. 2019. 'How Efficiency Shapes Human Language' *Trends in Cognitive Science* 23: 389-407.

Hansen, Nat. 2011. 'Color Adjectives and Radical Contextualism' *Linguistics and Philosophy* 34: 201-221.

Harris, Daniel. 2022. 'Semantics Without Semantic Content' *Mind* & *Language* 37: 304-328.

Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. PhD Thesis. University of Massachusetts, Amherst.

Hintikka, Jaakko. 1977. 'Quantifiers in Natural Language: Some Logical Problems II' *Linguistics and Philosophy* 1: 153-172.

Holton, Richard. 1997. 'Some Telling Examples: A Reply to Tsohatzidis' *Journal of Pragmatics* 28: 625-628.

Hesni, Samia. 2021. 'Generics as Instructions' *Synthese* 199: 12587-12602.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux.

Kaplan, David. 1989. 'Demonstratives', in J. Almog, J. Perry, and H. Wettstein, eds., *Themes from Kaplan*. Oxford University Press.

Knobe, Joshua et al. 2013. 'Dual Character Concepts and the Normative Dimension of Conceptual Representation' *Cognition* 127: 242-257.

Lee, Junhyo and Anthony Nguyen. 2021. 'What's Positive and Negative About Generics: A Constrained Indexical Approach' *Philosophical Studies* 179: 1739-1761.

Leslie, Sarah-Jane. 2007. 'Generics and the Structure of the Mind' *Philosophical Perspectives* 21: 375-403.

Leslie, Sarah-Jane. 2008. 'Generics: Cognition and Acquisition' *Philosophical Review* 117: 1-47.

Leslie, Sarah-Jane. 2012. 'Generics', in G. Russell and D. Graff Fara, eds., *The Routledge Companion to Philosophy of Language*. Routledge.

Leslie, Sarah-Jane. 2015*a*. 'Generics Oversimplified' *Noûs* 49: 29-54.

Leslie, Sarah-Jane. 2015*b*. '"Hillary Clinton is the Only Man in the Obama Administration": Dual Character Concepts, Generics, and Gender' *Analytic Philosophy* 56: 111-141.
Leslie, Sarah-Jane. 2017. 'The Original Sin of Cognition: Fear, Prejudice, and Generalization' *Journal of Philosophy* 114: 393-421.
Leslie, Sarah-Jane, Khemlani, Sangeet, and Glucksberg, Sam. 2011. 'Do All Ducks Lay Eggs? The Generic Overgeneralization Effect' *Journal of Memory and Language* 65: 15-31.
Nguyen, Anthony. 2020. 'The Radical Account of Bare Plural Generics' *Philosophical Studies* 177: 1303-1331.
Nickel, Bernhard. 2009. 'Generics and the Ways of the Normality' *Linguistics and Philosophy* 31: 629-648.
Nickel, Bernhard. 2010. '*Ceteris Paribus* Laws: Generics and Natural Kinds' *Philosophers' Imprint* 10: 1-25.
Nickel, Bernhard. 2012. 'Saying and Doing: The Role of Semantics in the Use of Generic Sentences' *Canadian Journal of Linguistics* 57: 289-302.
Nickel, Bernhard. 2016. *Between Logic and the World: An Integrated Theory of Generics*. Oxford University Press.
Plunkett, David, Sterken, Rachel Katherine, and Sundell, Timothy. 2023. 'Generics and Metalinguistic Negotiation' *Synthese* 201: 1-46.
Recanati, François. 2010. *Truth-Conditional Pragmatics*. Oxford University Press.
Roberts, Craige. 1996. 'Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics' *Semantics & Pragmatics* 5: 1-69.
van Rooij, Robert. MS. 'The Predictability Tree.'
Saul, Jennifer. 2017. 'Are Generics Especially Pernicious?' *Inquiry* February: 1-18
Sorensen, Roy. 2012. 'The Sorites and the Generic Overgeneralization Effect' *Analysis* 72: 444-449.
Sterken, Rachel Katherine. 2015*a*. 'Generics, Content and Cognitive Bias' *Analytic Philosophy* 56: 75-93.
Sterken, Rachel Katherine. 2015*b*. 'Leslie on Generics' *Philosophical Studies* 172: 2493-2512.
Sterken, Rachel Katherine. 2015*c*. 'Generics in Context' *Philosophers' Imprint* 15.
Sterken, Rachel Katherine. 2016. 'Generics, Covert Structure and Logical Form' *Mind & Language* 31: 503-529.
Sterken, Rachel Katherine. 2017. 'The Meaning of Generics' *Philosophy Compass* 12: 1-13.
Stokke, Andreas. 2013. 'Protagonist Projection' *Mind & Language* 28: 204-232.
Thakral, Ravi. 2018. 'Generics and Weak Necessity' *Inquiry* February: 1-28.
Yalcin, Seth. 2018. 'Semantics as Model-Based Science'. In D. Ball and B. Rabern, eds., *The Science of Meaning*. Oxford University Press.