Generics revisited: a simple theory of types

Exam number: B125648

Word count: 8,448[1]

MSc Taught Philosophy

The University of Edinburgh

2017/2018

[1] With the permission from the thesis supervisor and the postgraduate office

Contents

Abstract

In this paper I propose a novel treatment of genericity. While my diagnosis is that no present theory is satisfactory, I choose Liebesman's simple kind predication theory (2011) as an initial framework. I reject Liebesman's claim that generics are sentences about kinds and suggest that they are rather sentences about types showing the virtues of such view. While my ultimate goal is to show that the type analysis of generics is both natural and explanatorily superior to its rivals, I also propose both semantic and pragmatic analysis of generic sentences.

Keywords: *generics, genericity, simple kind theory, type/token distinction, structural universals*

The pervasive view of generics is that they are quantified sentences containing an implicit quantifier *gen*. An alternative to this position was proposed by David Liebesman (2011), according to whom *gen* does not exist and generics predicate properties directly to kinds. In this paper I will follow Liebesman in denying the existence of *gen*. My view diverges from Liebesman's in two respects: first, invoking a well-known *type/token* distinction I claim that generics should be analysed as sentences about types rather than kinds. Second, I disagree that the logical form of a generic sentence *As are B* is that of a kind signified by *A* having a property *B*. Making instrumental use of the theory of structural universals, I suggest that the type denoted by *A*, has a structural property *B'*, which, though dependent on *B*, is a different one.

Type analysis, on my view, has the following virtues: first, type/token distinction can be naturally applied to any general name, while not every referring phrase can be so easily treated as referring to a kind (consider *letter 'i'*, in *the letter 'i' is written with a dot*, which clearly signifies type, but doubtfully a kind). Second, I will argue that the common features of types explain the problematic features of generic sentences; kind-analysis is inferior in such respect. Finally, the type-view suggests a natural truth-criterion of generic sentences, which is neither present in Liebesman's paper, nor follows naturally from the simple kind view. However, given the space constraints, I focus on articulating the virtues of type-view rather than criticising the kind-view.

The structure of this paper is as follows: I will first discuss the phenomenon of genericity and the basic worries surrounding their analysis. I will then discuss the most prominent theories of genericity and show that none of them provide a satisfactory answer to those worries. Next, I

argue that the majority of the problematic features of generics are reducible to corresponding features of types; consequently, there are good grounds to think that generics are sentences about types. I will finally propose both rudimentary semantics and pragmatics of generics.

What are generics?

True generic sentences:

1.

a. Horses are mammals,
b. Dogs bark,
c. Ticks carry Lyme disease,
d. Dodos are extinct,
e. Buffalo form protective circles,
f. Frenchmen eat horsemeat,
g. Dobermans have floppy ears (context of biology),
h. Dobermans have pointed ears (context of breed standards),
i. Cats lick themselves,
j. A gentleman does not offend a lady,
k. The cheetah is the fastest mammal on earth,
l. Dutchmen are good sailors,
m. Kangaroos live in Australia
n. Men are taller than women

False generic sentences:

2.

a. Lawyers have odd insurance numbers (in a world w where by some accident rather than a rule all lawyers have odd insurance numbers)
b. Lions are female
c. Books are paperback
d. Dogs don't bark

Generics proved themselves highly resistant to a uniform analysis, and it is hard even to provide anything more than an ostensive definition. The name *generic*, stemming from Latin *genus* (meaning *kind*) suggests that generics are sentences about kinds. They indeed seem to universally use kind-names in subject-place rather than singular ones; however, it is unclear whether they are semantically about kinds (quantificational accounts of generics argue against such view). It is also clear, that not all sentences that use kind-names are generics, example being universal sentences. Thus, containing a kind-name is, at best, a necessary condition. Generics are usually expressed using bare plurals, but as shown by Leslie, Khemlani & Glucksberg (2011) and Leslie & Gelman (2012), some universal sentences are interpreted as generics when such interpretation is available (thus, one can speak about genericity as a semantic phenomenon, spread against the variety of syntactic forms). Also, as shown by Cohen & Erteschik-Shir (2002), many bare plurals have existential readings, and at least some of them have universal readings. Thus, the bare plural form is typical of generics, but constitutes neither necessary nor sufficient conditions for genericity. A more intuitive explanation of what generics are is given by Collins: 'What all these cases share is that under a normal, out-of-the-blue reading, their truth conditions concern not how things are for a given specific number of individuals <…> but how things are in some looser generic sense.' (2018: 35) I will not provide a stricter definition. For the purposes of this paper and given a vast literature on the topic I will suppose that it is clear what generics are.

Difficulties in theorising about generics

The following is the list of the more infamous properties of the generics.

1) **Generics vary in their quantificational power**. 1a requires not only that all factual horses are mammals, but also that all possible horses are, i.e., suggests a necessary universal quantification. However, others, such as 1b and 1c allow for exceptions. Notoriously, *lions have four legs* seems to remain true even if after some accident (think combat, madman cutting legs, etc.) all lions had less than 4 legs, showing that generics might remain true even if no individuals of the factual world correspond to what is claimed.
2) **Generics differ in predicative force**. While 1a and 1b say what can be predicated to individuals, predicate in 1d can be applied only to kinds. Examples such as 1i, also the

so-called *donkey-sentences* (sentences where a variable is intuitively bound but where standard rules of syntax prohibit binding of the variable, such as in *donkey owners beat their donkeys*) which might express generic facts, suggest that in some cases predicates involved can only be predicated to individuals.

3) **Generics differ in their expressive mode**. Some generics express well-established truths (1a), others express striking properties typical to a group (1f), others still express norms (1j).
4) **Generics are context sensitive**. As apparent from 1g and 1h, at least some generics seem to have different interpretations and different truth values in different contexts.

Combined with the already mentioned absence of unique syntactic form, this makes generics immune to a simple analysis. We arrive at what is known as the paradox of acquisition (due to Leslie, 2008):

5) **The paradox of acquisition**. Though being a property of theories of genericity rather than generics themselves, it is an important observation that though generics are among the easiest linguistic devices to master (being competently used by children who still struggle with explicitly quantified sentences), they are incredibly resistant to theorizing.

Semantic theories of generics

Any interesting account of genericity targets to explain at least some of these properties. Ideally, one simple and powerful theory would account for first four, and then, given its simplicity, would answer the paradox of acquisition.

There are two main models of theorizing about the meaning of generics: quantificational and referential. According to the quantificational model, a generic sentence *As are B* quantify over the individuals described by A. Quantification might (or might not) have modal import. Generally, the logical form of the sentence *As are B* would be expressed by $gen\ x\ [A(x)][B(x)]$. Proponents of the second model suggest that generics employ reference directly to kinds, and their logical form is something like this: $\lambda x[B(x)](A')$, where A' is a referring term, and B is a predicate term. For the sake of brevity, I outline only selected arguments for and against the both

camps. The more extensive overview might be found in Sterken (2016, 2017) and Nickel (2016; especially Chapters 2 and 3).

The quantificational account is the majority position; David Liebesman even calls it the 'consensus view' (2011: 409). According to Liebesman, the view is mainly grounded on the following two observations: 1) adding quantificational adverbs such as *typically* or *generally* does not alter the meaning of certain generic sentences (consider *dogs bark* and *typically, dogs bark*), thus, one may think that such generics contain implicit, unpronounced quantifier, similar to *typically* or *generally*. 2) Subject term in a generic sentence is a common noun, which is usually analysed as a predicate with an argument place which requires saturation; and while the saturating element is not present in the surface structure, it might plausibly be implicit.

Apart from these two, another strong argument concerns subject dependent predicates. Consider *Cats lick themselves*. Simple kind predication theories predict that their logical form should be $\lambda x[x\ licks\ x](Cat)$, and hence their meaning would be rendered to that of *cats lick cats*, which is an undesired result.

Simple kind predication can be supported by the following arguments. First, it is striking that though according to quantificational theories, there exists an (implicit) quantifier *gen*, no known language contains a pronounced version of it, though generics are available in all known spoken languages. This is a strong evidence that *gen* does not exist, and thus, that generic sentences are unquantified. Second, one might consider a well known *de dicto/de re* distinction, which comes into play in intensional contexts, and notice that generics behave as unquantified sentences. Consider:

a) Marry believes that all celebrities do drugs

b) Mary believes that John does drugs

c) Mary believes that celebrities do drugs

Now a) has two available readings: Mary might believe that some certain people, namely, *x, y, …* (the list contains all factual celebrities) do drugs (the *de re* reading). Or she might not have beliefs about particular people, but rather think that whoever is a celebrity does drugs (the *de dicto* reading). However, b), which contains no quantification, has no corresponding ambiguity.

Nickel (2016: 36) notices that c) similarly has no ambiguity involved, which supports non-quantificational view.

Finally, there are sentences, such as 1d or 1k, which suggest that no quantification is involved, for the predicates seem to be true of kinds rather than individuals.

Much of the inspiration for both camps comes from the sentences of different predicative force; and naturally, much of the debate goes around these differences. In this respect, neither part seems to demonstrate clear superiority, and their success is measured by their strengths in other departments.

Quantificational theories come in three main varieties: probabilistic, normality and meta-semantic. Probabilistic theories are best represented by Ariel Cohen (2012). According to Cohen, generics are sentences about probabilities; in particular, *As are B* say that As are probable to be B. More specifically, Cohen believes that the operator *gen* is modal, and that it quantifies within the worlds that share their history with the actual one and do not undergo any unexpected changes, i.e., are uniform. The operator's force is probabilistic: if x is A holds, then within the relevant worlds $P(B(x)|A(x)) > 0.5$. One obvious counterexample to such theory is 2b: in the worlds that share their history to the actual one, and which are not to change in any striking way, the probability that a randomly chosen lion is female is >0.5. Additional counterevidence comes from 1c, 1f and 1k, showing that Cohen's theory both misdiagnoses some true generics as false, and some false generics as true.

Normality theories are exemplified by Nickel's (2016) account, according to which *gen* quantifies over close possible worlds restricted to all normal individuals, where normality is understood as normality in a certain respect. E.g., *As are Bs* means that there is a way of being a normal A (with respect to B), such that all normal As in all relevant worlds are Bs. This allows Nickel to account not only for differences in quantificational power (1), but also, at least partially, for differences in expressive mode (3), and context sensitivity (4). Regarding expressive mode, normality claims are intuitively interpreted as suggestive of what the actual world should look like (the word *normal* itself has normative force). Context sensitivity is accounted for various ways of being normal; thus, apparent incongruency of 1g and 1h can be explained away by noticing

that there are two ways of being a normal Doberman with respect to ear shape: to either have pointed, or floppy ears.

However, as noticed by Sterken (2017), this account misdiagnoses *books are paperbacks* as true. Also, it diagnoses 1k (or, for a bare plural case, the alternative *cheetahs are the fastest mammals on earth*) as false. Thus, Nickel's account once again, provides both false positives and false negatives. Yet another counterexample to Nickel's theory comes from double genericity, example being 1m. Boundaries of normal men and women height intersect, and thus there is a sufficiently close possible world in which at least some normal men (with respect to their height) are shorter than some normal women, falsely predicting the sentence to be false.

A prominent example of a metasemantic theory is one due to Sterken (2015). According to her, semantics of generics is simple: their logical form is $gen\ x\ [A(x)][B(x)]$ and this is all there is to say. *Gen,* however, is context-sensitive (which explains the context-sensitivity of some, or according to Sterken, almost all generic sentences). In particular, it is an indexical quantifier. Thus, the variations of quantificational power are to be answered by a metasemantic account of how *gen* picks up different quantifiers in different contexts. Sterken's account thus answers why there are differences in predicative power (1), and it can be easily employed to explain the differences in expressive mode (3), for there are quantifiers which suggest universal prevalence over possible worlds; there are weak ones, quantifying over small groups of individuals within the actual world; and there are some which suggest normative reading. Indexicality of *gen* automatically explains context sensitivity (4). While all the hard work is shifted to metasemantics, the semantics of generics is simple, which answers the paradox of acquisition (5). Thus, by its explanatory power Sterken's account is superior to others described in this paper.

However, the account has its troubles. Apart from the general difficulties of quantificational accounts, the very idea that *gen* is indexical is, though certainly tempting, rather dubious. First, typically indexicals cannot be evaluated as true or false without knowing the context of utterance (consider *She is rich, Today is Tuesday*, *The Beatles once had a concert in that city*). They do not have a preferred reading in the absence of context. However, majority of the generics have natural readings which do not require any context at all. In fact, all the examples listed at the beginning of this paper, with an exception of 1h and 1g, can be naturally evaluated as true or

false without any given context. Second – and this is a methodological weakness of Sterken's account – is that it is underdeveloped to the point that it is a scaffold for future theories of generics rather than a theory itself. Sterken does not propose any specific metasemantics for *gen.* One is left with an impression that the complex properties of generics are explained away, but they are simply moved to metasemantics, which is yet to be analysed.

Simple kind predication is exemplified by Liebesman's account. According to Liebesman, A in *As are B* refers to a kind, and, roughly, this is the end of the semantic story. Logical form of the generic sentence is as simple as *F(k)*, where a property denoted by *F* is predicated to a kind *k*. According to Liebesman, the complex properties of generics are to be explained outside their semantics; that is, he does not think that a semantic theory of genericity should account for their properties listed above. His proposed semantics being simple explains away the paradox of acquisition.

Simple kind predication revisited: outline

My proposal is closest to that of a simple kind predication. I will take Liebesman's (2011) account as my starting point. As noticed above, simple kind predication theories struggle in analysing essentially individual predicates. Sterken suggests that Liebesman can account for them by claiming that they are existential rather than generic (e.g., *cats lick themselves* meaning simply *some cats lick themselves*), and that the resulting weakness of such analysis can be remedied by employing Gricean pragmatics. Leslie (2015) argues that Liebesman might employ type-shifting operator to type-shift down a kind term when needed; e.g., *Tigers are striped*, which on Liebesman's analysis is as simple as Striped(Tigers), can be analysed as [(⇓(tigers)(x)][striped(x)], 'where ⇓ is a type shifting operator taking a kind term to a predicate that applies to all and only member of that kind.' (Leslie, 2015: 30) I will follow neither of the suggestions.

Ultimately, Leslie diagnoses Liebesman's account as explanatorily inadequate. I agree that its major methodological vice is that instead of explaining the complexity of generics' properties it simply shifts the burden of explanation to metaphysics. As noticed by Collins (2018), Liebesman does a good job at criticising the existence of *gen*, but his account is not really *simple*, for it does not simplify the phenomenon of genericity in any way. His account is methodologically similar

Sterken's (2015): it seems to be quite unproblematic, but at the price of a proportionate explanatory shallowness of the theory.

Any adequate proposal, according to Leslie, must account for the following phenomena:

1) The non-equivalence of sentences such as *tigers are kinds* and *Panthera tigris is a kind*
2) The underlying complexity in logical form that is revealed by sentences such as *cats lick themselves* and *politicians think that they can outsmart their opponents*
3) Donkey anaphora in generics, as in 'children who grow a new tooth like to show it off'
4) Weak crossover effect in the generic context [e.g., it *His mother loves John* is fine, but *their mothers love boys* is not]
5) The phenomenon of indefinite singular generics whose indefinite singular subjects cannot refer to kinds [e.g., *A mosquito carries West Nile virus*] (Leslie, 2015: 22)

In what follows I will propose my account of a simple type view, arguing that it answers the above worries.

What are types?

How many letters are there in the word *letter*? Either there are four, namely *l, e, t, r*, or there are six (*e* and *t* occur twice). When we choose the former answer, we speak about letters as types. We also speak about letters as types when we say that *a* is the first letter of the Latin alphabet. When we choose the latter, we speak about instances, or tokens. We also speak about letters as tokens when we say that a certain letter *b* is blurry or written in pink ink. Letter tokens are physical entities, they might be imprinted on page, displayed on a screen etc.; types are non-physical. The distinction is by no means restricted to letters. One might speak about a horse as a type, when one says *the horse has a tail*, or a horse as a token, when one says *this horse is old*. Wetzel and Zalta (2014) suggest that when 'we read that *the Spirit Bear is a rare white bear that lives in rain forests along the British Columbia coast*, we know that no particular bear is *rare*, but rather a type of bear. <…> It is even more evident that a type is being referred to when it is claimed that *all men carry the same Y chromosome* <…>'. One should, however, remain careful with one's intuitive understanding of type/token distinction. A mistake to be avoided is that of identifying *tokens* with *occurrences.* As noticed by Wetzel and Zalta, (ibid.) if one takes a famous line from Gertrude's Stein *Sacred Emily*, 'rose is a rose is a rose is a rose', and counts the occurrences of the word *rose*

in it, one can do so by counting occurrences both in a sentence type, or in a sentence token. If one speaks about 'the famous sentence from *Sacred Emily*', one might not mean a particular copy of the poem; and thus speak about poem-type, sentence-type, and the occurrences of *rose* in it, meaning word *rose* as a type, and each of the occurrences of the word *rose* as type-occurrences. We thus might legitimately speak about both type-occurrences and token-occurrences.

I will take that the distinction is clear and will now list certain facts regarding type/token distinction, which will be useful in the further discussion of generics.

First, types are related to properties in different way than their tokens, which stems from their ontological differences. The prevalent view is that types are abstract, and there is a good reason to think so: it seems that types are not spatiotemporal – e.g., there is no particular location where *the* letter *a* exists; types do not cease to exist once their tokens do; they seem to be mostly causally inert; and these are typical traits of abstract entities. But if so, then the usual way of speaking about types is somehow ambiguous. While it is true to say *the basketball is round*, the type of a basketball itself is a non-physical entity and has no physical shape: types cannot be round. And while it is true that *Roman centuria consists of 100 soldiers*, this can be literally true only of tokens of centuria; for physical entities – soldiers – can make up only another physical entity, and not an abstract one.

It will not do to say that the examples just given are not indeed about types, and that genuine predication to types does not involve the ambiguity in question. For consider '*a' is the first letter of the Latin alphabet.* No particular token is the first letter of the Latin alphabet, thus what is said must be said of a type. However, ambiguity remains: type *a* is not a letter, on a par with type *horse* not being an animal (but rather an abstract entity), and what is said is neither literally true of any token, nor of the type. Thus, while the sentences in question are true sentences about types, their meaning is non-literal. We may conclude that types do not *have* the properties predicated (in a sense tokens have them), and that the way of speaking we employ in uttering sentences about types is not literal. After all, one can follow Quine in interpreting types as sets of properties (1987: 218); it is then clear that the relation between a set and its elements is different from that of an object and its properties. One helpful way to emphasize the predicative difference is to say that tokens *have* properties, and that types *entail* them, though both facts are expressed

using the same linguistic means (i.e., both *this horse has a tail* and *the horse has a tail* uses the same mechanism of predicating a property, though first expresses an individual having a property, and second expresses a type entailing a property).

Second, tokens can have properties other than those inherited from types. A token of the letter *a*, written in invisible ink, has a property of being invisible. However, letter *a*, as a type, does not entail the property of invisibility. A particular horse may lack a tail, even though 'horse', as a type, entails having a tail. Tokens might have altogether new properties (think of a horse wearing a hat), or they can have properties that overwrite inherited ones (think, again, *taillessness*). Since objects, once thought of as the tokens of a certain type, are to resemble that type, a natural thought is that in some ideal scenario all the tokens would be exemplifications of the type. However, tokens are physical objects that undergo changes, some of the inherited properties are lost, and some new are acquired, and thus either type/token analysis is not suitable for a physical world, or some divergence of tokens from their type is to be allowed. A token *x* of a type *X* might undergo so many changes that it can no longer be called a token of *X*. I will not go into a debate on how many of the properties of a token of some type T can be overwritten for an entity to remain plausibly a token of that type. This was discussed to some extent by Saul Kripke in his famous lectures *Naming and Necessity*, and in the inspired literature; but for my needs it is enough to conclude that token properties, over and above the inherited ones (and even conflicting with inherited ones) are to be expected.

Third, type-talk forces a *de dicto* reading, which, in turn, affects modality. Take *John collects coins.* Assuming that the sentence is true, it is hard to articulate which coins John has collected, or how many of them. It indeed seems irrelevant; it might be the case that John got robbed yesterday, and he has no coins in his collection. However, if he does not intend to quit his hobby, it is still true to say that he collects coins. It seems intuitive to me that the target of John's activity is not some individuals *per se*, but some things (factual or possible) just because they are coins; and thus the target is not only factual, but also possible individuals (i.e., whatever is, or might be, a coin). The natural reading of *coins* in *John collects coins* is then a *de dicto* one. Also, the target of John's activity is clearly a type *coin,* rather than some certain coin tokens. The aforementioned ambiguity occurs: John collects physical coins, rather than abstract objects, so what is said is not that John

collects the coin type; but it is neither something about any particular set of coins. How many, or which coins John gets to collect is entirely accidental.

Finally, types can have varieties. E.g., the letter *a,* as a type, can be either the written sign, or the sound one makes. Which one is used is normally recognized from context. E.g., if we speak about majuscule and minuscule, we clearly speak about letters as written signs; and when we speak about silent letters (in, say, French language) we speak about letters as sounds. It would be misleading to say that the written letter *x* in the French word *deux,* and the silent letter *x* in the French word *deux* are different letters: the sentence *the letter x in the French word 'deux' is written, but not pronounced* seems true, but it could hardly be analysed as such if one chose to postulate two different types, namely orthographic and phonological letter *x.* It is thus more plausible to treat such cases as involving varieties of the same type rather than different types. Which variety is at use is established by context; thus, some context sensitivity in type-sentences is to be expected.

Genericity problems revisited

Let us review the features of generics listed at the beginning of this paper in light of what has been just said. My goal here is to show that the type analysis of generic sentences allows to reduce the features of generics to those of types.

Varying quantificational power. Tokens do not instantiate all the properties of types, or only those; hence we get differences in quantificational power. Type sentences do not quantify over individuals at all, and the question *how many* is irrelevant. However, metaphysically tokens tend to exemplify types, and the exceptions stem from the physical. It is easy to change colour, loose a leg or two; but this is hard to attain immortality; and this comes into play when comparing the prevalence of the properties in *horses have four legs* and *men are mortal*. Also, type-sentences involve the *de dicto* reading in virtue of what they are modally stronger than universal sentences: while not quantifying, they still speak about tokens in both factual and possible worlds. Consequently, the properties predicated, with the varying level of exceptions, are to be expected in the possible worlds. Now if the Kripkean revolution taught us something, it is that some properties cannot be lost without affecting the very identity of the individual; and thus some of

the properties occur without exceptions throughout the possible worlds. The issue involves properties of individuals rather than properties of types, and is thus not a problem of genericity.

These points answer what Cohen (2012) calls *the paradox of generics*: generics seem to be both stronger and weaker than universal sentences. The suggested answer is this: they are quantificationally weaker (for they allow exceptions) but modally stronger (for their scope overarches possible worlds).

Differences in predicative force. Types have properties over and above those of tokens, and some of them are type-bound. Speaking about the properties had, *all* of them differ from those that tokens have: generics are not to be read literally. However, there are interpretational differences: the sentences that use token-level predicates have natural interpretation as saying something directly of tokens, while those which involve type-bound predicates do not: e.g., *dogs bark* is naturally (but, according to my proposal, misleadingly) interpretable as saying about certain dogs that they bark; but *oaks are widespread* has no interpretation as a statement about particular oaks. Thus the apparent difference of predicative force is that of ambiguity involved; but the predicative force is uniform – that of predicating a certain structural property directly to a type.

Now types are tied more closely to tokens than kinds to their members; for a type is more of an exemplary, uninstantiated case of a token. Properties that types entail are those which can be instantiated by tokens.

Finally, when we have sentences such as *dodos are extinct*, the type-bound property *extinct* has a straightforward interpretation.

Differences in expressive mode. The issue is twofold: semantic and pragmatic. Semantically, generics express truths about types, and metaphysically tokens are to resemble types. This in turn adds some natural normativity to type-talk: if one is to be a legitimate token of some kind, she is to resemble that type. E.g., if *men don't cry* is true, one – at least if he wants to be legitimately called a man, i.e., to be a legitimate token of type *man* – should not cry.

Pragmatically, as I will suggest at the end of this paper, generics communicate expectations (this stems naturally from the fact that future tokens are expected to continue to

resemble types); and expectations can be understood both factually and normatively. One can expect it to rain tomorrow; or one can expect you to never offend a lady.

Generics are context sensitive. This stems from the metaphysics of types. As noticed in the discussion above, types might have varieties. One way to articulate this is by saying that there might be different sorts of the same thing (where *thing* is a type: e.g., letter A, as a type; Doberman, as a type), i.e., different sorts of the same type. Depending on the context, the relevant variation is picked, and the meaning is adjusted. Thus, for example, we may name two sorts of Doberman (as a type): that of a biological kind; and that of a breed standard.

These points help answer Leslie's **paradox of acquisition:** generics are hard to theorize about, since they are metaphysically complicated. They are, however, easy to master, since, as I will argue later, they are both semantically and pragmatically simple.

Semantics of generics

My claim is that generics are simple statements about types; a sentence *As are B* says that type A entails property B. It is now time to provide a definition of the entailment relation, which, as mentioned above, amounts to providing a theory of genericity. I employed it to capture the following difficulty: the linguistic means of type-talk suggest that some property is predicated to type (consider *The rhino has a horn made of keratin*, which is a true sentence about rhino type), but it cannot be predicated to type (for types are abstract entities with no physical parts, and consequently no parts made of keratin). It is thus clear that the underlying logical form is not straightforwardly represented.

Armstrong (1986) fought the same difficulty by defending the theory of *structural universals*. His worry was the following: methane molecules, which instantiate the universal *methane*, are composed in a certain way of one carbon atom and four hydrogen atoms. Now it is true to say about the universal *methane* that it is composed of carbon and hydrogen atoms; but these atoms cannot be physical, for they constitute a universal. Armstrong suggests that the universal *methane* is composed of universals *carbon* and *hydrogen* in a fashion similar to that of methane instances being composed of carbon and hydrogen instances. *Methane* is then a structural universal, which is structurally similar to real methane, but lacks its physical features. Following Armstrong one can then say that instead of having physical properties, types have

structural properties corresponding to those of their tokens. I.e., instead of having a property B, type A has some other property, B', determined by B. Since types on such view are structural entities, we might say that the type A which entails B, has a structural counterpart of a property B. Making use of such proposal I will denote a structural counterpart of a property B by B_T. Thus, for example, Red_T, or type-redness, is a property which a type has, if it entails redness. I am, however, making no metaphysical commitments here; as shown by Bigelow and Pargetter (1989), one might leave the ontological questions of universals aside while making instrumental use of the structural analysis.

It is helpful to notice next that type-bound predicates, such as *exists* or *widespread,* do not describe structural properties; i.e., they do not correspond to physical properties of tokens. There is no ambiguity involved; what is meant is a property which is had, not entailed, by type. I am thus inclined to claim that such sentences are not true generics after all; however, for a time being I leave this question open.

Up to now the semantics of generics seem to be really simple: I suggest that for a type-bound property P and type t their logical form is as simple as Liebesman suggests, i.e., *P(t).* In case of non-type-bound properties the logical form is that of entailment, which can be reduced to that of type having an equivalent structural property; i.e., for a non-type-bound property P' and type t it is $P'_T(t)$, contrary to Liebesman's simple proposal of $P'(t)$.

But the remaining worry is this: what does it mean that a type has a structural property? The challenge is partially epistemological: we plausibly obtain knowledge about types via their tokens. The properties we are most familiar with are token-level properties, and the relation between our normal objects and their properties is that of having, rather than entailing. We cannot move to some metaphysical plane without explaining how we are able to competently speak about such plane, and thus if we are not to leave the theory explanatorily inadequate we must describe the rules of reduction, i.e.: what does it mean, on the level of tokens, that a type entails a token-level property.

My suggestion is this: having a structural property P_T amounts to all of the tokens inheriting a corresponding property P. Inheritance here is to be understood causally: if a token inherits property from a type, that token will be within a scope of a causal mechanism bringing

about *P*, and will likely have *P*, unless something more particular than being a token of that type causes it to lack such property, i.e., causes an exception. The amount of exceptions, depending partially on the nature of the property in question, might vary drastically. Thus, the suggestion amounts to saying that if *As are B* is true, then all tokens of A are within a scope of a certain causal mechanism bringing about B. Since causality is not to be understood as necessarily bringing about the effect (consider *smoking causes cancer*), but rather as increasing the likelihood of the effect, exceptions might have two sources: 1) causal mechanism failing to generate the effect 2) some other cause altering the effect.

It seems that if a token has a property which is not inherited from type, there must be some external cause for it, i.e., some other causal mechanism than that involved in being a token of that type. Conversely, if a property is inherited from type no external cause is needed: if it is true that dogs have four legs, and there is some individual dog who has four legs, there does not need to be any further cause from him simply being a dog (and thus exercising certain causal mechanism) that he has four legs. Similarly, if we see a dog wearing a hat, and it is false that dogs wear hats (i.e., that the type *dog* entails wearing a hat), then it is not simply in virtue of it being a dog that it wears a hat; there must be some other reason – some external cause – of it wearing a hat (i.e., someone dressing it up).

But now, surely, there might be cases where a token of A has a property P in virtue of some external cause, even though *As are P* is true. E.g., there might be an albino raven painted in black. It might even be the case that there are no As which are P (e.g., a possible world in which a madman left all dogs legless, in which it is still intuitively true that *Dogs have four legs*), or that all As that are P are such in virtue of some external causes (as in 2a). Thus if we are to speak about tokens demonstrating inherited properties we clearly need modal import. This is also supported by the already discussed forced *de dicto* reading: type-talk involves modalities, and what is said about type is, usually, not bound to a certain world.

These considerations lead to the following criterion:

A has a structural property P_T (alternatively: type *A* entails property *P) if and only if there could exist a token of A which has a property P simply via being a token of a type A*

A straightforward objection occurs at this point: there are some distant possible worlds, namely, those in which the negation of a certain true generic *As are B* is true, with respect to which the criterion makes a false prediction (e.g., a distant possible world in which dogs evolved to have three legs renders *dogs have three legs* true, since there could, if the world had such unexpected turn, that there was a dog which would have three legs just via being a dog). Thus, clearly, what is needed is quantification over close possibilities. I suggest the following option. Imagine a scenario in which all dogs are left with three legs, though dogs generally have four legs. Now there is a possible world in which dogs have four legs and at least one dog has four legs. And this world is certainly closer than that in which dogs generally have three legs. Thus our search through possible worlds should end at the closest suitable world. The updated criterion is then this:

A has a structural property P_T (alternatively: type *A* entails property *P) if and only if there exists a token of A in a sufficiently close possible world which has a property P simply via being a token of a type A (i.e., there exists no closer possible world in which some token of A lacks a property P (e.g., has an incompatible property P*) simply via being a token of a type A).*

For the sake of simplicity in the following examples I will assume that we stop at a sufficiently close possible world.

Analysis of the examples

1a is a straightforward case. It does seem that there might be some individual horse which is a mammal without any external cause, i.e., in virtue of simply being a horse. Indeed, all of the horses, including all possible ones (within history-sharing worlds) are.

1b involves a slight ambiguity, stemming from a phenomenon of habituality. Habituals are sentences such as *Mary smokes after work,* which describe a recurring behaviour. A generic *dogs bark* uses the predicate *bark* in a habitual sense. And for almost every dog to habitually bark one does not need any external causes (i.e., someone teaching it how to produce a barking sound, genetically engineering it, etc.). Thus, there is, or at least might be, a dog which (habitually) barks without any external causes, i.e., simply via being a token of a type *dog*.

An apparent counterexample to my proposal comes from 1c. Surely, some particular tick will not carry a Lyme disease just because it is a tick; it needs to bite an infected animal first. There

is something about being a tick – mosquitos do not carry Lyme disease even if they bite an infected animal – thus plausibly there is some causal mechanism involved, but not a sufficient one. Thus, it seems that being a tick is a component cause, at best. I think, contrary to these intuitions, that the issue is pragmatic at heart. Take *The bride stands on the left*. This generic sentence seems true, but its natural interpretation involves a relevant context, i.e., that of a wedding ceremony. It is false that the bride always stands on the left when one speaks about the wedding photos, which is another possible context of interpretation. Thus, the proposition in mind is expressed by *the bride always stands in the left during the wedding ceremony*, and the property entailed by the type *the bride* is that of standing in the left during the wedding ceremony. 1c, I believe, is similar in this aspect; for it targets only the relevant situations; i.e., those in which a tick has bitten an infected animal. Thus, we do not need to update the criterion; we must rather notice that some clues of what the exact property is entailed by type might come from the context.

1e is slightly problematic in yet another way: the predicate is not a token-level one. The entailment of the property *form protective circles* does not translate straightforwardly into a simple token-level predicate, i.e., to a property had by token; but it is still easy to make sense of. Let us denote a relevant individual-level predicate by *B*, which is an abbreviation of the entailment relation *E*(*x, form protective circles*). Now a token of buffalo, to demonstrate a relevant behaviour (that of joining a protective circle when it is being formed) in a habitual way, i.e., to have a property B, does not need to be forced by some external causes; the mechanism is inherited from him being a buffalo, and thus it is true that there might exist a buffalo which has a property B simply via being a buffalo.

1f, as noticed by multiple authors (e.g., Sterken, 2015), involves context sensitivity, and thus, as discussed before, suggests that variation in type is involved. The context in which it is true to say that Frenchmen eat horsemeat suggests some historically traditional Frenchman as a type; and the causal mechanism which affects the whole type (and in virtue of that all of its tokens) is probably that of a tradition. It is then clear that there might be a token of such type, which will habitually eat horsemeat not by some external causes (e.g., being forced to), but via him being a Frenchman (in that sense). This might be articulated by the following conversation:

- *Did you know that John eats horsemeat? What is wrong with him?*

- *Ah, he is just French.*

The answer implies that being French in some sense explains why John eats horsemeat, and thus suggests that there is a causal mechanism involved between being French and habitually eating horsemeat.

1g is once again straightforward; if Dobermans indeed have floppy ears, there will be tokens of Dobermans which will have floppy ears just in virtue of them being Dobermans. 1h, on the proposed account, involves a variety of type *Doberman*; in particular, that of the breed standard. A possible token of the breed-standard-Doberman will indeed have pointed ears not because of any external causes; the fact that its ears will be cut is not *external*, for it is a part of being breed-standard-Doberman. Thus, its ears will be cut just because it is a breed-standard-Doberman, and not because, say, it was mistaken with Schnauzer (which would indeed be an external cause).

1i is unproblematic. The property of self-licking is predicated to tokens, while what is predicated to a type is some structural counterpart, *self-licking*$_T$. Since type relates to tokens in a different way than kind to its members, we do not get an apparent absurdity; what we get, intuitively, is that a cat type has a structural property of *licking itself*, and by exemplification its tokens lick themselves (rather than cat type).

The account gives a desirable prediction for 1j. There might be a gentleman – an exemplary one – which does not offend a lady just because he is a gentleman.

1k, as was noticed above, is a tricky case; but the test proposed does fine. Indeed, if the cheetah is the fastest mammal on earth, it is possible that there will be a cheetah, which is the fastest mammal on earth, not because it was trained/genetically engineered, etc., but just because it is cheetah.

1l is tricky in that I am unsure whether this sentence is even correct. Moreover, if it is, then it is clear that some context-sensitivity is involved here, and the type in question is a specific variety of a type *Dutchman*, probably in a similar sense as a Frenchman in 1f. But, I believe, if the sentence is true, then there might exist a certain Dutchman who is a good sailor just because he is a Dutchman. If all of the Dutchmen who are good sailors (including possible ones) are good

because they have attended a particular school, or just by some accident, i.e., if there are always external causes involved, then it will be false that Dutchmen are good sailors.

1m is again simple: there might exist a particular Kangaroo which lives in Australia simply in virtue of it being a Kangaroo. Indeed, majority of the kangaroos which live in Australia live there without any external causes; and the causal mechanism which is involved is probably some mixture of evolution and the continent of Australia being isolated from the rest of the world by water.

1n involves genericity twice; for it is unclear which particular men are to be taller than which particular women. My proposal need a slight update to account for it. Intuitively, what is being said is that man, as a type, is taller than woman, as a type; thus, what is expressed is a certain relation. This, naturally, renders the suggestion to the following: there could possibly exist such man and such woman, that that man is taller than that woman simply because he is a man, and she is a woman, i.e., not because of any external causes.

Let us now turn to the false cases.

2a is false, since there is no possible lawyer who has an odd insurance number just via him being a lawyer; i.e., there is no implicit causal mechanism, a rule, etc., and if he is to get an odd insurance number, it needs some external cause (e.g., coincidence, him asking to have it, etc.)

2b is straightforwardly false, for no lion is female just via it being lion. External cause is a particular mechanism, e.g., having a certain chromosome. Now surely, each lion has a sex just via it being a lion. But which sex in particular it will have is unclear, unless more particular mechanisms are involved.

2c is false, for no paperbacks are paperbacks just in virtue of them being books in sufficiently close worlds.

Regarding 2d, there are breeds of dog which do not bark, e.g., Borzois. The situation here is similar to that of 2b: surely, a dogs have a breeds just in virtue of them being dogs (and thus *dogs have breeds* is true), but what particular breed it has is an external question. Thus, a possible dog, who does not bark because it is a Borzoi, does so not just because it is a dog; but rather because it is a Borzoi, rendering the sentence false.

Genericity and pragmatics

What is said above has useful implications for the pragmatics of the generic sentences. I have suggested that generics are sentences about types; and a natural interpretation is that types are helpful generalizations over individuals. Since types concern not only factual, but also possible tokens, these generalizations concern the world as it might be. Thus, our knowledge about types give us reasonable and helpful expectations about the world, and this, I suggest, is the pragmatic function of generic sentences: they are used to communicate what to expect – generally, or in a particular context. It might be helpful to expect a tick to carry Lyme disease; a certain person, in some situation, to have weird eating habits (e.g., eat horsemeat); a dog to bark; or to expect that a gentleman will not offend you, if you are a lady. Kangaroos are expected to be encountered in Australia, and you are not expected to encounter a Dodo.

This then shows that pragmatically they are far more primitive than Leslie (2007) suggests; they do not involve separate categories, and not access some implicit cognitive capacities. They are to communicate simply what is to be expected of the tokens of a particular kind.

But then it is only natural that generics have different expressive modes: for *expectation* clearly invokes a normative force. The fact that a gentleman is expected to not offend a lady might be true both in virtue of some mechanism, which forces gentlemen not to offend ladies; or it might be the expectation itself which makes the sentence true.

Implications

The major virtue of the type-view over simple kind analysis is its explanatory power: as suggested above, the metaphysical properties of types and the linguistic features of type-talk together account for the fundamental problematic features of generics. In addition to that, the type-view can easily answer what remains of the list of Leslie's problems. The non-equivalence of sentences such as *tigers are kinds* and *Panthera tigris is a kind* is explained straightforwardly: *tigers are kinds* would imply *kinds*$_T$(tiger), meaning that the type *tiger* entails the property *is a kind*, reducing to particular token tigers being kinds, which is false. However, *Panthera tigris is a kind* is not a generic sentence, but a true individual-level sentence, meaning that *Panthera tigris* has a property *is a kind*. Finally, indefinite singular generics whose indefinite singular subjects cannot refer to kinds by the suggested treatment simply do not refer to kinds; they can, however, refer to types.

I believe that the type-analysis is a move in a correct direction, and making such move was indeed the aim of this paper. However, the criterion provided here is a very rough first attempt based on the general insights coming from the type-view; a more thorough investigation is yet to be done.

Bibliography

1. Armstrong, D. (1986). In Defence of Structural Universals. *Australasian Journal of Philosophy*, 64: 85-88.
2. Bigelow, J., & Pargetter, R. (1989). A theory of structural universals. *Australasian Journal of Philosophy*, *67*(1), 1-11
3. Cohen, A. (2012). Generics as modals. *Recherches linguistiques de Vincennes*, (41), 63-82.
4. Cohen, A., & Erteschik-Shir, N. (2002). Topic, focus, and the interpretation of bare plurals. *Natural Language Semantics*, *10*(2), 125-165.
5. Collins, J. (2018). Genericity sans Gen. *Mind & Language*, *33*(1), 34-64.
6. Leslie, S. (2007). Generics and the structure of the mind. *Philosophical Perspectives*, 21(1), 375-403.
7. Leslie, S. (2008). Generics: cognition and acquisition. *Philosophical Review*, 117(1), 1-47.
8. Leslie, S. J. (2015). Generics oversimplified. *Noûs*, 49(1), 28-54.
9. Leslie, S. J., & Gelman, S. A. (2012). Quantified statements are recalled as generics: Evidence from preschool children and adults. *Cognitive Psychology*, 64(3), 186-214.
10. Leslie, S. J., Khemlani, S., & Glucksberg, S. (2011). Do all ducks lay eggs? The generic overgeneralization effect. *Journal of Memory and Languag*e, 65(1), 15-31.
11. Liebesman, D. (2011). Simple generics. *Noûs*, 45(3), 409-442.
12. Nickel, (2016). Between logic and the world: An integrated theory of generics. Oxford, UK: Oxford University Press.
13. Quine, W. V. (1987). *Quiddities: An Intermittently Philosophical Dictionary*. Cambridge, MA: Harvard University Press.
14. Sterken, R. K. (2015). Generics in context. *Philosopher's Imprint*, 15(21), 1–30.
15. Sterken, R. K. (2016). Generics, covert structure and logical form. *Mind & Language*, *31*(5), 503-529.
16. Sterken, R. K. (2017). The meaning of generics. *Philosophy Compass*, *12*

17. Wetzel, L., Zalta, E. N. (2014). Types and Tokens. *The Stanford Encyclopedia of Philosophy (2014 Spring Edition)* URL = <https://plato.stanford.edu/archives/spr2014/entries/types-tokens/>.