

Generic Sentences: Representativeness or causality?

Robert VAN ROOIJ— *ILLC, University of Amsterdam*

Katrin SCHULZ — *ILLC, University of Amsterdam*

Abstract. Many generic sentences express stable inductive generalizations. Stable inductive generalizations are typically true for a causal reason. In this paper we investigate to what extent this is also the case for the generalizations expressed by generic sentences. More in particular, we discuss the possibility that many generic sentences of the form ‘*ks* have feature *e*’ are true because (members of) kind *k* have the causal power to ‘produce’ feature *e*. We will argue that such an analysis is quite close to a probabilistic based analysis of generic sentences according to which ‘relatively many’ *ks* have feature *e*, and that, in fact, this latter type of analysis can be ‘grounded’ in terms of causal powers. We will argue, moreover, that the causal power analysis is sometimes preferred to a correlation-based analysis, because it takes the causal structure that gives rise to the probabilistic data into account. Unfortunately, there are problems for the causal power analysis too, and we will discuss them as well.

Keywords: Generic sentences, causality, probability.

1. Introduction

The proposal that we will discuss in this paper is that many generic statements, like (1a)-(1c) should be given a causal analysis.

- (1) a. Tigers are striped.
- b. Birds fly.
- c. Birds lay eggs.

This paper is structured as follows: in the following section we will briefly motivate a recently proposed frequency-based descriptive analysis according to which a generic sentence of the form ‘*ks* are *e*’ express inductive generalizations. In this section we will also discuss a conceptual problem for this frequency-based analysis: the fact that the analysis seems too extensional. In section 3 we will provide a causal explanation for the descriptive analysis making use of some natural independence assumptions, and thereby solving the above mentioned conceptual problem that the frequency-based analysis of section 2 is too extensional. In section 4 we will show that once the independence assumptions of our causal derivation are given up, a causal analysis will, perhaps, give rise to improved empirical predictions. We will argue that giving up these assumptions is sometimes very natural. In section 5 we will discuss another important distinction between a purely probabilistic account, and the causal one suggested in this paper: the resulting asymmetry due to causality and the need for an additional evidential reading. Section 6 concludes the paper.

2. A probabilistic analysis of generics and its problems

Generic sentences come in very different sorts. Consider (2a), (2b) and (2c).

- (2) a. Tigers are striped.
- b. Mosquitoes carry the West Nile virus.
- c. Wolves attack people.

We take (2a) to be true, because the vast majority of tigers have stripes. But we take (2b) and (2c) to be true as well, even though less than 1% of mosquitoes carry the virus and the vast majority of wolves never attack people. Most accounts of generics, if they don't stipulate an ambiguity, start from examples like (2a) and then try to develop a convincing story for examples like (2b) and (2c) from here. In our previous analysis (van Rooij and Schulz, to appear), in contrast, we took examples like (2b) and (2c) as points of departure and then generalized the analysis to account for more standard examples as well, in the hope that it would lead to a more uniform analysis.

What is the natural analysis of examples like (2b)? We take this to be that:

1. it is **typical** for mosquitoes that they carry the West Nile virus, and
2. this is highly relevant information, because of the **impact** of being bitten by a mosquito when it carries the West Nile virus.

We take it that it is intuitively quite clear when one feature has a significantly higher **impact** than another. This is normally the case when the first feature gives rise to a more negative emotional reaction than the latter. We realize, however, that from here to come up from a quantitative measure of 'impact' is a long way to go, and we have not much to offer.¹

As for **typicality**, it is obviously not required for *e* to be a typical feature for *ks* that *all ks* have feature *e*. Although almost all tigers are striped, there exist albino tigers as well, which are not striped. And although '(be able to) fly' is a typical feature for birds, we all know that penguins don't have this feature. The same examples show that *e* can be typical for *k* although not only *ks* have feature *e*: cow and cats, too, can be striped, and bats fly as well. So we need a weaker notion of typicality. We take it that *distinctiveness* matters for typicality, and thus for generics. This can be illustrated by the contrast between (3a), which is intuitively true, versus (3b), which is false.

- (3) a. Lions have manes.
b. *Lions are male.

One might think that (3b) is false because only 50%, if at all, of the lions are male, which cannot be enough for a generic to be true. But that, clearly, cannot be the reason: the only lions that have manes are male lions. Thus, not even 50% of the lions have manes. Still, (3a) is, intuitively, true. The conclusion seems obvious: (3b) is true, because it is *distinctive* for lions to have manes, where the notion of distinctiveness shouldn't be too strong. On a weaker analysis of 'being distinctive', one demands only that *in comparison with* other larger animals, many males have manes. Similarly, for (2b) to be true it is at least required that compared to other insects, many mosquitoes carry the West Nile virus. To account for this comparative analysis, one could make use of either a qualitative or a quantitative analysis. But because we want to incorporate the importance of the second condition, *impact*, within an analysis of 'relatively many', it is almost mandatory to provide a *quantitative* analysis of distinctiveness.^{2,3} If typicality

¹However, we do have one suggestion: looking at news items. What is typically being reported in news items are things or events that we feel have a big impact, even if they are rather uncommon.

²Of course, these considerations are well-known to users of decision- and game theory. who have to combine uncertainty with utility.

³This argument won't have any force if one takes generic sentences to be *ambiguous* between majority-generics

reduces to distinctiveness and if we have such a quantitative analysis of distinctiveness, plus a quantitative measure of impact, we can define a measure of a generic sentence of the form ‘*ks* are *e*’ as distinctiveness of *e* from *k*, $Distinctiveness(e, k)$, \times $Impact(e)$. We could call this measure the *representativeness* of *e* for *k*. Because we will argue later that typicality cannot always be reduced to distinctiveness, the representativeness of *e* for *k*, $Repr(e, k)$, should be defined more generally as

- $Repr(e, k) =_{df} Typicality(e, k) \times Impact(e)$.

Then we can say that the generic sentence ‘*ks* are *e*’ is true, or appropriate, if and only if the *representativeness* of *e* for *k*, $Repr(e, k)$, is high:

- ‘*ks* are *e*’ is true if and only if $Repr(e, k)$ is high.

Before we concentrate on the more general notion of typicality, let us first discuss various potential measures of distinctiveness. To provide a quantitative analysis of what it means that feature *e* is distinctive for group *k*, i.e., that relatively many *ks* have feature *e*, there are many options open. On one natural analysis, it holds that relatively many *ks* have feature *e* if and only if the relative frequency of *ks* that are *e* is higher than the relative frequency of alternatives of *k* that are *e*. If we measure relative frequency by probability function *P*, this can be captured by the condition that $P(e|k)$ — i.e., the conditional probability of having feature *e* given that one is a member of group or kind *k*—is higher than $P(e|\cup Alt(k))$, where $Alt(k)$ denotes the (contextually given) alternatives to group *k*, and $\cup Alt(k)$ thus denotes the set of members of any of those alternatives. For readability, we will from now on abbreviate $\cup Alt(k)$ by $\neg k$. Thus, relatively many *ks* are *e* iff $P(e|k) - P(e|\neg k) > 0$. In psychology, the measure $P(e|k) - P(e|\neg k)$ is called ‘contingency’ and denoted by ΔP_k^e . This notion plays an important role in the theory of associative learning (cf. Schanks, 1995), and it is well-known that $\Delta P_k^e > 0$ if and only if $P(e|k) > P(e)$, the standard notion of *relevance*.⁴ It should be noted, however, that $P(e|k) - P(e|\neg k)$ does not behave monotone increasing with respect to $P(e|k) - P(e)$.⁵ So the choice between these two measures makes a difference for predictions. Notice that if we use contingency to model distinctiveness, and if also typicality reduces to it, it is predicted that the generic ‘*ks* are *e*’ is true, or acceptable, if and only if $[P(e|k) - P(e|\neg k)] \times Impact(e)$ is high. This, in turn, is high iff $P(e|k) \times Impact(e) \gg P(e|\neg k) \times Impact(e)$, if ‘ \gg ’ means ‘highly above’. If we abbreviate $P(e|k) \times Impact(e)$ by $EV(e|k)$, the expected value of *e* given *k*, this means that the generic ‘*ks* are *e*’ is true, or acceptable, iff $EV(e|k) \gg EV(e|\neg k)$. If we use standard relevance, the generic would be true iff $EV(e|k) \gg EV(e)$. For features with $Impact(e) = 1$, these two equalities hold iff $P(e|k) \gg P(e|\neg k)$ and $P(e|k) \gg P(e)$, respectively, meaning that a small difference between $P(e|k)$ and $P(e|\neg k)$ (or $P(e)$) is not enough to make the generic true.

Although contingency is a natural measure to determine ‘distinctiveness’, in van Rooij and Schulz (to appear) we proposed that it should be measured by a slight variant of Shep’s (1958)

like (2a), on the one hand, and ‘striking’ generics like (2b) and (2c), on the other. In fact, Leslie (2008) proposed such an ‘ambiguity’-analysis. But we don’t see any empirical evidence in favor of such an ambiguity analysis, and we thus take it to be obvious that a uniform analysis is preferred. We will see that what Leslie calls ‘majority’-generics fall out as a special case of our uniform analysis.

⁴Cohen (1999) proposed that a generic sentence of form ‘*ks* are *e*’ is true on its relative reading iff $P(e|k) > P(e)$, if we limit the ‘domain’ of the probability function to $k \cup \cup Alt(k)$.

⁵In fact, $P(e|k) - P(e) = P(\neg k) \times [P(e|k) - P(e|\neg k)]$.

notion of ‘relative difference’, $\Delta^{**}P_k^e = \frac{\alpha P(e|k) - (1-\alpha)P(e|\neg k)}{\alpha - (1-\alpha)P(e|\neg k)}$, with $\alpha = \frac{P(k)}{P(k)+P(\neg k)}$. Notice that in case $P(k) = \frac{1}{2}$, it will also be the case that $\alpha = \frac{1}{2}$ and that $\Delta^{**}P_k^e$ comes down to Shep’s notion $\Delta^*P_k^e = \frac{P(e|k) - P(e|\neg k)}{1 - P(e|\neg k)}$, while in case $P(\neg k) = 0$ —i.e. when $\bigcup Alt(k) = \emptyset$ — α ends up being 1 and $\Delta^{**}P_k^e$ comes down to $P(e|k)$. We will assume that the tokens of the alternative kinds are chosen such that $P(\bigcup Alt(k)) = P(\neg k) = P(k)$, in case $Alt(k) \neq \emptyset$. Thus, $\alpha \in \{\frac{1}{2}, 1\}$. In sum:

$$\bullet \textit{Typicality}(e, k) =_{df} \frac{\alpha P(e|k) - (1-\alpha)P(e|\neg k)}{\alpha - (1-\alpha)P(e|\neg k)} = \Delta^{**}P_k^e \quad \text{with } \alpha = \frac{P(k)}{P(k)+P(\neg k)}.$$

Three arguments were given for this choice:

(i) Suppose that the vast majority of members of $\bigcup Alt(k)$ are of kind k' and that $P(e|k')$ is slightly higher than $P(e|k)$. If we don’t control for the number of tokens of alternative kinds, or types, we take into account, ‘ ks are e ’ will be predicted to be false, even if for most $k'' \in Alt(k)$ $P(e|k) \gg P(e|k'')$. But that seems wrong. One way to predict correctly would be to count not *all tokens* of the alternatives types, but rather equally many tokens of each alternative type such that we look at as many tokens if we look at the tokens of all these types together as there are tokens of k . Thus, it is important that we control for the number of tokens of alternative kinds, and the demand that $P(\bigcup Alt(k)) = P(\neg k) = P(k)$ is a special case of this.

(ii) in case $P(e|k) = 1$ and $P(e|\neg k) \neq 1$, the generic sentence seems to be perfect, whatever the value of $P(e|\neg k)$ is. In contrast to the standard notion of relevance, this comes out by using our measure of typicality for both values of α .

(iii) in case e is an uncommon feature, i.e. when $P(e|\neg k)$, or $P(e)$, is low, the difference between $P(e|k)$ and $P(e|\neg k)$ — $P(e|k) - P(e|\neg k)$ —should be larger for the generic to be true or appropriate than when $P(e|\neg k)$ is high, if $\alpha = \frac{1}{2}$.⁶

From (ii) and (iii) it follows that for distinctiveness of e for k , the conditional probability of e given k , $P(e|k)$, counts for more than $P(e|\neg k)$. And this seems required. Consider, on the one hand, the uncommon feature ‘having 3 legs’. Although there are (presumably) relatively more dogs with three legs than there are other animals with three legs, this doesn’t mean that the generic ‘Dogs have three legs’ is true (cf. Leslie, 2008). If a more common feature is used, on the other hand, an equally small difference between $P(e|k)$ and $P(e|\neg k)$ can make the difference between truth and falsity of the generic sentence, if the generic is used to contrast k from other kinds.

In summary, the following analysis of generic sentences of the form ‘ ks are e ’ was proposed:

- ‘ ks are e ’ is true if and only if $Repr(e, k)$ is high.
- $Repr(e, k) =_{df} \Delta^{**}P_k^e \times Impact(e)$.
- $\Delta^{**}P_k^e =_{df} \frac{\alpha P(e|k) - (1-\alpha)P(e|\neg k)}{\alpha - (1-\alpha)P(e|\neg k)}$ with $\alpha = \frac{P(k)}{P(k)+P(\neg k)}$.

It should be clear how examples like (2a)-(2c) can be accounted for on this proposal: (2a) is true, or appropriate, because being striped is distinctive for tigers, whereas (2b) is true because (i) more mosquitos than other types of insects carry the West Nile virus, and (ii) carrying this

⁶For instance, in case $P(e|\neg k) = 0.9$, the value of $\frac{P(e|k) - P(e|\neg k)}{1 - P(e|\neg k)}$ is $10 \times [P(e|k) - P(e|\neg k)]$, while if $P(e|\neg k) \approx 0$, the value of $\frac{P(e|k) - P(e|\neg k)}{1 - P(e|\neg k)}$ is just $P(e|k) - P(e|\neg k)$, so 10 times smaller.

dangerous virus has a high impact. In van Rooij and Schulz (to appear) it is argued that a wide variety of generics can be accounted for using the above analysis, especially if (i) we make use of the context-dependence of which alternatives are relevant, and (ii) we assume that it is not just relative frequency that counts, but rather *stable* relative frequencies: it is not only that the measure $P(e|k) - P(e|\neg k)$ should be high, but this measure should remain high when conditioned on relevant backgrounds.⁷ Moreover, in van Rooij and Schulz (to appear) it is argued that a high value of $Repr(e, k)$ gives rise —partly due to Tversky & Kahneman’s Heuristics and Biases approach — to the (perhaps false) impression that $P(e|k)$ is high, thereby accounting for the general intuition that generics like ‘*ks* are *e*’ are true just in case $P(e|k)$ is high. This is not the place to defend this view, and the bulk part of this paper is written on the (hypothetical?) idea that the above analysis of generics is accepted to be roughly correct.

One obvious objection to the above descriptive analysis in terms of (stable) frequencies should be mentioned, though: $\Delta^{**}P_k^e$ by itself cannot account for the ‘intensional component’ of generic sentences showing in their ‘non-accidental’ understanding. Even if actually (by chance) all ten children of Mr. X are girls, the generic ‘Children of Mr. X are girls’ still seems false or inappropriate.⁸ The sentence only seems appropriate if being a child of Mr. X somehow *explains why* one is a girl. In this paper we will explore to what extent we can explain the meaning of generic sentences in terms of *inherent dispositions* or *causal powers*. Even though such dispositions were philosophically suspect in much of the 20th century, we take such an exploration as a worthwhile enterprise, because it seems to be in accordance with many people’s intuition. Moreover, by adopting a causal stance, the non-accidental understanding of generics can, arguably, be explained as well.

3. Causal readings of generics

3.1. Causal explanation of correlations

The theory of generics in terms of the measure $\Delta^{**}P_k^e$ is very Humean, built on frequency data and probabilistic dependencies and the way we learn from those. Many linguists and philosophers feel that there must be something more: something hidden underlying these actual dependencies that **explains** them. A most natural explanation is a **causal** one: the probabilistic dependencies exists *in virtue of* objective kinds which have causal powers, capacities or dispositions.⁹ Indeed, traditionally philosophers have assumed that the natural world is objectively divided into kinds, which have essences, a view that has gained popularity in the 20th century again due to the work of Kripke and Putnam. A closely associated modern view that has gained popularity recently has it that causal powers (Harre and Madden, 1975), capacities (Cartwright, 1989) or dispositions (Shoemaker, 1980; Bird, 2007) are the truth-makers of laws

⁷The notion of stability is required to think of $P(e|k) - P(e|\neg k)$ as helping to account for inductive generalizations, and does the work that Cohen (1999) argues his condition of ‘homogeneity’ should do. It is by concentrating ourselves on probabilities that are *stable* under conditionalization by various conditions that generics like ‘Bees are sterile’, ‘Israeli live along the coast’ and ‘People are over three years old’ are predicted to be bad, or false, although in each case the majority of the ‘kind’ has the relevant feature. For an analysis of stability that we favor, see Skyrms (1980).

⁸To account for such cases, the ‘unbounded’ character of generics, Cohen (1999) makes use of limiting relative frequencies. The solution we will propose in the following sections will be different, but based on a similar intuition.

⁹It seems no accident that (general) causal statements typically are of generic form: ‘Sparks cause fires’, ‘Asbestos causes cancer’.

and other generalities.¹⁰ Tversky and Kahneman's (1980) observed experimentally that the idea is natural: ordinary people are inclined to a causal interpretation of correlations.

While probabilistic (in)dependencies are symmetric, causal relations are not. But neither are generic sentences. Such sentences of the form 'ks are e' are, by their very nature, stated in an asymmetric way: first the noun *k*, then feature *e*. This naturally gives rise to the expectation that objects of type *k* are associated with features of type *e* because the former has the power to cause the latter. Where the goal of van Rooij and Schulz (to appear) was to develop a semantic analysis of generic sentences that is descriptively adequate, the goal of this paper is to investigate to what extent this theory can be explained by basing it on an analysis of (perhaps unobservable) causal powers. In a sense, the answer to this question is quite clear: Shep's notion of relative difference closely corresponds to Good's (1961) measure of 'causal support': $\log \frac{P(\neg e|\neg k)}{P(\neg e|k)}$. In fact, Good's notion is ordinal equivalent to Shep's notion in the sense that $\Delta^* P_k^e > \Delta^* P_{k^*}^{e^*}$ iff $\log \frac{P(\neg e|\neg k)}{P(\neg e|k)} > \log \frac{P(\neg e^*|\neg k^*)}{P(\neg e^*|k^*)}$ for all e, e^*, k and k^* . This is very interesting. In the end, though, also Good's notion is just a frequency measure. What we would like to find is a 'deeper' foundation of our measure. In a sense, this is what Good provides as well, for he provides an axiomatization of his notion of causal support. But we think that the causal foundation that we will give is more natural, and fundamental.

We don't want to claim that a causal analysis can account for all types of generics. Generics like 'People born in 1990 reach the age of 40 in the year 2030' and 'Bishops move diagonally' (in chess) are most naturally not treated in a causal way. Linguists (Lawler, 1973; Greenberg, 2003) also make a difference between generics formulated in terms of bare plurals (BP) ('Dogs bark'), on the one hand, and generics stated in terms of indefinite singular (IS) noun phrases ('A dog barks'), and found that IS generics have a more limited felicity, and suggested that in contrast to a BP generic, for an IS generic to be felicitous, there has to exist a 'principle connection' between subject noun and predicate attributed to it. Perhaps this means that only IS generics should be given a causal analysis. Perhaps. But we do think that for many, if not most, BP generics causality could play an important role as well. The purpose of this paper is not to defend the strong view that all generics should be analyzed causally. Instead, our purpose is more modest: to explore the possibility of a causal power analysis of BP generics.¹¹

As part of this, we want to clarify what, if any, advantage(s) such a causal power analysis might provide. These advantages could be of a conceptual and an empirical nature. As for the former, if all that is gained by a causal analysis of e.g., 'Aspirin relieves headaches' is that the observed frequency of relieved headaches is said to be due to the Aspirin's unobservable capacity to relieve headache, nothing is won. For a causal analysis to be useful more insights should be gained, for instance in the internal structure of the cause. But a causal analysis will be useful here as well, as shown by the recent abundance of papers on *mediation* (Preacher and Kelley, 2011; Pearl, 2014): the use of causal models to explain not only *why* something happened, but also *how* it happened. Scientists are not only interested to learn *that* Aspirin relieves headaches, they are also interested in the *mechanism* by which it does so. In this paper we will make use of the recent insights of causal mediation analyses that make a difference between *direct* and *indirect* causal effects. In the next section we will show that under certain circum-

¹⁰ According to Strawson (1989), even Hume himself believed in causal powers.

¹¹ We leave a discussion of IS generics to another paper.

stances a causal interpretation gives rise to different, and arguably *more adequate predictions* than an extensional theory making use of $\Delta^{**}P_k^e$, especially if we concentrate on *direct effects*. But first we will show in this section that under natural assumptions a causal analysis *explains* the predictions made by using $\Delta^{**}P_k^e$.

3.2. A causal derivation of $\Delta^{**}P_k^e$

For our causal explanation of the measure $\Delta^{**}P_k^e$ we follow Cheng (1997) and assume that objects of type k have unobservable **causal powers** to produce features of type e . We will denote this unobservable causal power by p_{ke} . It is the probability with which k produces e when k is present in the absence of any alternative cause. This is different from $P(e|k)$. The latter is the relative frequency of e in the presence of k . We will denote by u the (unobserved) alternative potential cause of e (or perhaps the union of alternative potential causes of e), and by p_{ue} and $P(e|u)$ the causal power of u to produce e and the conditional probability of e given u , respectively. We will assume (i) that e does not occur without a cause and that k and u are the only potential causes of e (or better that u is the union of all other potential causes of e other than k), (ii) that p_{ke} is independent of p_{ue} , and (iii) that p_{ke} and p_{ue} are independent of $P(k)$ and $P(u)$, respectively, where independence of p_{ke} on $P(k)$ means that the probability that k occurs and produces e is the same as $P(k) \times p_{ke}$. The latter independence assumptions are crucial: by making them we can explain the *stability* and (relative) context-independence of generic statements.

Now we are going to derive p_{ke} , the causal power of k to produce e , following Cheng (1997).¹² To do so, we will first define $P(e)$ assuming that e does not occur without a cause and that there are only two potential causes, k and u (recall that $P(k \vee u) = P(k) + P(u) - P(k \wedge u)$):

$$(4) \quad P(e) = P(k) \times p_{ke} + P(u) \times p_{ue} - P(k \wedge u) \times p_{ke} \times p_{ue}.$$

Thus, the probability that e , conditional on k and $\neg u$ is

$$(5) \quad P(e|k, \neg u) = p_{ke} = \text{the causal power of } k \text{ to generate } e.$$

Because the probability that e given k in the absence of all other causes u does not depend on the frequency of these other causes, it does not depend on the base rate of e . As such it can be thought of as an inherent property of k , and can be extrapolated across contexts in which base rates differ.

One problem with this notion is that it depends on u , which is unobservable, or many times cannot be identified. Thus, it still remains unclear how to estimate the causal power of k to produce e . It turns out that things are clearer if we *assume that k and u are, or are believed to be, independent of each other*. Assuming independence of k and u , $P(e)$ becomes

$$(6) \quad P(e) = P(k) \times p_{ke} + P(u) \times p_{ue} - P(k) \times P(u) \times p_{ke} \times p_{ue}.$$

As in section 2, ΔP_k^e is going to be defined in terms of conditional probabilities:

$$(7) \quad \Delta P_k^e = P(e|k) - P(e|\neg k).$$

¹²Glymour (2001) calls Cheng's derivation 'a brilliant piece of mathematical metaphysics'. Cheng (1997) also discusses an analysis of *preventive* causes. We won't deal with this in this paper.

The relevant conditional probabilities are now derived as follows:

$$(8) \quad \begin{aligned} P(e|k) &= p_{ke} + (P(u|k) \times p_{ue}) - p_{ke} \times P(u|k) \times p_{ue}. \\ P(e|\neg k) &= P(u|\neg k) \times p_{ue} \quad (\text{derived from (4), because } P(k|\neg k) = 0). \end{aligned}$$

As a result, ΔP_k^e comes down to

$$(9) \quad \begin{aligned} \Delta P_k^e &= p_{ke} + (P(u|k) \times p_{ue}) - (p_{ke} \times P(u|k) \times p_{ue}) - (P(u|\neg k) \times p_{ue}) \\ &= [1 - (P(u|k) \times p_{ue})] \times p_{ke} + [P(u|k) - P(u|\neg k)] \times p_{ue}. \end{aligned}$$

From this last formula we can derive p_{ke} as follows:

$$(10) \quad p_{ke} = \frac{\Delta P_k^e - [P(u|k) - P(u|\neg k)] \times p_{ue}}{1 - P(u|k) \times p_{ue}}.$$

From (10) we can see that ΔP_k^e gives a good approximation of causal power in case (i) u is independent of k (meaning that $P(u|k) - P(u|\neg k) = 0$), and (ii) $p_{ue} \times P(u|k)$ is low. Obviously, in case k is the only potential direct cause of e , i.e., when $p_{ue} = 0$, it holds that $p_{ke} = \Delta P_k^e$. Because in those cases $P(e|\neg k) = 0$, it even follows that $p_{ke} = P(e|k)$.

Our above derivation shows that to determine p_{ke} in case events or features of type e might have more causes, we have to know the causal power of p_{ue} , which is equally unobservable as p_{ke} . You might wonder what we have learned from the above derivation for such circumstances. It turns out, however, that p_{ke} can be estimated in terms of observable frequencies after all, because we assumed that $P(k)$ and $P(u)$ are independent of each other. On this assumption it follows that $P(u|k) = P(u) = P(u|\neg k)$ and that (10) comes down to

$$(11) \quad p_{ke} = \frac{\Delta P_k^e}{1 - P(u|k) \times p_{ue}}.$$

Because of our latter independence assumption, it follows as well that $P(u|k) \times p_{ue} = P(u) \times p_{ue} = P(e|\neg k)$. This is because $P(u) \times p_{ue}$ is the probability that e occurs and is produced by u . Now, $P(e|\neg k)$ estimates $P(u) \times p_{ue}$ because k occurs independently of u , and, in the absence of k , only u produces e . It follows that p_{ke} can be defined in terms of observable frequencies as follows:

$$(12) \quad p_{ke} = \frac{\Delta P_k^e}{1 - P(e|\neg k)}.$$

But this is exactly the same as $\Delta^* P_k^e$, the measure in terms of which we have stated the truth conditions of generic sentences in section 2! Thus, in case we assume that a generic sentence of the form ‘Objects of type k have feature e ’ is true because objects of type k cause, or produce, features of type e , we derive exactly the semantics we have proposed in the first place (if $\alpha = \frac{1}{2}$). It follows that as far as our descriptive analysis of generics in terms of $\Delta^* P_k^e$ was correct, what we have provided in this section is a **causal explanation**, or grounding, of this descriptive analysis.

The above derivation causally motivated Shep’s notion of ‘relative difference’. But that notion is a special case of $\Delta^{**} P_k^e$ in case $\alpha = \frac{1}{2}$. We have seen above that in case $\alpha = 1$, what should come out is that $\Delta^{**} P_k^e$ comes down to $P(e|k)$. Does a causal analysis motivate this as well? It does! We have seen in section 2 that $\alpha = 1$ just in case $Alt(k) = \emptyset$. In our causal derivation, this means that there is no alternative potential cause of e , i.e., that k is the *only potential cause* of e . What happens to our above derivation in that case? To see this, notice that in that case $P(e)$ can be determined as follows:


$$(13) \quad P(e) = P(k) \times p_{ke}.$$

As a result, $P(e|k)$ reduces to p_{ke} . Thus, $p_{ke} = P(e|k)$ in case k is the only potential cause of e , just like $\Delta^{**}P_k^e$ came down to $P(e|k)$ in case $Alt(k) = \emptyset$. Thus, our earlier measure $\Delta^{**}P_k^e$ could be motivated by our causal powers view both when $\alpha = \frac{1}{2}$ and when $\alpha = 1$.

How do these causal powers account for generic sentences? This is easiest to see for generics involving homogenous substances, like ‘Sugar dissolves in water’ and ‘Metal conducts electricity’. Intuitively, these are true, because of the causal power of sugar and metal to generate the observable manifestations that come with the relevant predicates. Similarly, ‘Tigers are striped’ is true, on a causal account, because of what it is to be a tiger. But sometimes the power description should be relativized. For instance, ‘Ducks lay eggs’ is true, although only the female chickens do so. Intuitively, it is not the causal power of ‘being a duck’ in general that makes this generic true. Rather, it is the causal power of being a *female* duck. But this comes out naturally. Cohen (1999) argued that the ‘domain’ of the probability function should be limited to individuals that make at least one of the natural alternatives of the predicate term true. In our example, it is natural to assume that $Alt(lay\ eggs) = \{Lay\ eggs, give\ birth\ live\}$. Because $\cup Alt(lay\ eggs) \approx Female$, this means that we should only consider female ducks. This should be done as well for the estimation of causal power. Doing so, it will be the case that the causal power of female ducks to lay eggs is high, which gives rise to the correct prediction that the generic ‘Ducks lay eggs’ is true. It is also clear how our analysis can account for ‘striking’ generics like (2b) and (2c): instead of demanding that $\Delta^*P_k^e \times Impact(e)$ is high, one now demands that $p_{ke} \times Impact(e)$ is high, which normally comes down to the same.

4. Giving up independence of the potential causes

In the previous section we assumed with Cheng (1997) that e had two potential causes, k and u , and that these causes were independent of each other: $P(u|k) = P(u|\neg k) = P(u)$. As noted by Glymour (2001), by adopting this assumption, Cheng assumed implicitly a specific type of causal structure: that what via Pearl (1988) is known as a ‘Noisy OR-gate’.



Thus, as noted by Glymour (2001), the models that Cheng uses to calculate how we can estimate causal powers are in fact special cases of structural causal models as developed by Pearl (2000) and Spirtes et al. (2000).¹³ In general, the potential causes of a variable don’t have to

¹³These models, in turn, are also generalizations of structural analyses used in biology, genetics and epidemiology. The standard structural models in these social sciences are *linear models*, models which consists of a system of linear equations. In such a linear equation, the relationship between two variables, X and Y , can be given by something like $Y = \alpha + \beta X + \epsilon$, where α and β are constants, ϵ an error term, and Y and X the dependent and independent variable, respectively. Although the values of α and β (the path coefficient) can be determined by standard regression analysis, what is special about the structural interpretation of such linear models is that X is seen as a *cause* of Y : a change of the value of X causes a change of the value of Y , and not the other way around.

be independent of each other. It could be, for instance, that values of (variables) K and U both have causal influence on the value of E , but that K also has a causal influence on U , or the other way around. It could also be, of course, that there is a common cause of K and U .

We will see below (with Glymour, 2001)¹⁴ that also in such situations, the causal power of X to influence Y can sometimes be estimated from frequency data, at least if we keep in mind the causal structure that generated these data.

If independence is only a useful, but sometimes incorrect, heuristics to determine probabilities, it raises the question what happens if we give up this independence assumption? Quantitatively speaking, there are two possibilities: $P(u|k) > P(u|\neg k)$ and $P(u|k) < P(u|\neg k)$. Already by looking at the general definition of p_{ke} :

$$(14) \quad p_{ke} = \frac{\Delta P_k^e - [P(u|k) - P(u|\neg k)] \times p_{ue}}{1 - P(u|k) \times p_{ue}},$$

we can immediately observe the following:

1. If $P(u|k) < P(u|\neg k)$, then $\Delta^* P_k^e$ **underestimates** p_{ke} .
2. If $P(u|k) > P(u|\neg k)$, then $\Delta^* P_k^e$ **overestimates** p_{ke} .

Thus, although giving up on independence doesn't allow us anymore to determine p_{ke} in terms of observed frequencies alone (because we now also need to know p_{ue} , $P(u|k)$ and $P(u|\neg k)$), giving up independence still potentially gives rise to interesting empirical consequences. In the following subsections we will look at both cases, and see that they give rise to interesting new predictions.

4.1. $\Delta^* P_k^e$ underestimates p_{ke} .

First, we will look at the most extreme case where $P(u|k) < P(u|\neg k)$, namely where u and k are **incompatible**. Notice that in that case $P(u|k) = 0$. The relevant conditional probabilities are then derived from (4) as follows: $P(e) = P(k) \times p_{ke} + P(u) \times p_{ue}$. From this we derive immediately that $P(e|k) = p_{ke}$, because $P(u|k) = 0$.

Thus, we see that in case k and u are incompatible, the causal power of k to produce e is the same as the conditional probability $P(e|k)$, just as was the case if k is the only cause of e . Perhaps this can explain the intuition people have that the acceptability of a generic sentence of the form ' ks are e ' goes with its conditional probability $P(e|k)$. Thus, although under natural independence conditions $p_{ke} = \Delta^* P_k^e$, this is no longer the case once k and u are not taken to be probabilistically independent.

On this interpretation, the value of β carries causal information. The causal models of Pearl and Spirtes et al. are a generalization of the above, because for them the equations don't have to be of linear form.

¹⁴Pearl (2000) has an alternative derivation of $\Delta^* P_k^e$, what he calls 'the probability of causal sufficiency', PS. Pearl derives PS — our $\Delta^* P_k^e$ — from the measure $P(e_k|\neg k, \neg e)$, the probability of e after intervention with k when you are in a state where k and e are false. In this derivation Pearl doesn't use the assumption that there is a statistically independent alternative cause, u , that may produce e . He substitutes this assumption with an assumption of monotonicity: that k never prevents e . Pearl doesn't make use of causal powers in the derivation of $\Delta^* P_k^e$, but uses an assumption of causality, or intervention, as primitive instead. One can think of Pearl's PS as a generalization of Cheng's causal power as well, because applicable in more situations than Cheng's notion.

Are there good examples of generic statements where k and u (the union of alternative causes of feature e) are incompatible, or where k is taken to be the only cause of e ? This depends very much on what one takes the alternative causes to be. Take any generic of the form ‘ ks are e ’. Let us assume that $P(e|k)$ is high. We have argued in section 2 that this is not always enough to make the generic true. But now suppose that ‘ k ’ denotes a kind of animal (e.g., ‘horse’) and that e is a feature like ‘having a heart’. If one makes the Aristotelian assumption that x is a member of a kind if and only if x has the *essence* of that kind, then it is natural that we take the alternative causes of (having feature) e to be (essences of) other kinds of animals. Thus, $u = \bigcup Alt(k)$, with k incompatible with u . If for the analysis of generics we adopt the measure $\Delta^*P_k^e$ (with k denoting horses and e denoting creatures with a heart), the generic ‘Horses have a heart’ is most likely counted as false, simply because $P(e|k) = P(e|\neg k) = P(e|\bigcup Alt(k))$, and thus $P(e|k) - P(e|\neg k) = 0$, meaning that also $\Delta^*P_k^e = 0 = \Delta^*P_k^e$, if $\alpha = \frac{1}{2}$. Thus, on a correlation-based analysis, the generic is predicted to be false if $\alpha = \frac{1}{2}$.¹⁵ On a causal power view, however, the sentence is predicted to be true, because now $p_{ke} = P(e|k) \approx 1$. Of course, that $p_{ke} = P(e|k)$ was due to the assumption that k and u (the union of alternative causes of feature e) are incompatible. Perhaps this view only makes sense once one makes the highly controversial Aristotelian assumption that it is the essences of kinds that have causal powers. But controversial as this assumption might be, psychologists like (Keil, 1989; Gelman, 2003) and others have argued that both children and adults tend to have essentialist beliefs about a substantial number of categories.¹⁶

4.2. $\Delta^*P_k^e$ overestimates p_{ke} .


An analysis of generic sentences of the form ‘ ks are e ’ in terms of causal powers is perhaps most natural when k denotes a *homogenous natural kind*. The reason is that if two objects are of the same homogenous natural kind, they are very *similar*. It is this similarity that allows us to *inductively infer* unobservable features of this object of natural kind k from observed features of other objects of natural kind k . But this inductive inference crucially relies on a close similarity between the members of natural kind k . For natural or artificial kinds, or groups, that allow for a larger variation, an analysis in terms of causal powers might seem, perhaps, less natural. We want to argue that a causal analysis might be insightful, after all. The reason is that a causal analysis can take into account the causal structure *behind* the observed frequencies, and can make a distinction between *direct* and *indirect* causal powers, and effects.

If u , like k , is a generative cause of e , p_{ke} will be lower than $\Delta^*P_k^e$ in the following three causal

¹⁵Of course, $\Delta^*P_k^e$ comes down to $P(e|k)$ if $\alpha = 1$.

¹⁶Danks (2014) represents concepts as graphical-model-based probability distributions (see also Sloman, 2005). He shows that all the most prominent models of concepts (the theory-based, the prototype-based, and the exemplar-based) can be modeled by such distributions. An exemplar-based model of a concept, for instance, according to which the connection between an individual d and a concept C should be based on the similarity between d and each of the exemplars of C , can be represented by a probability function over features, such that every pair of features are associated with one another, but that these associations are all due to an *unobserved common cause*. (Danks (2014) shows how to directly translate in both directions between an exemplar-based concepts (making use of similarities between the members) and a graphical-based probability function with a common cause structure. Arguably, this is just as well the correct representation of a probabilistic version of a more traditional essence-based model of concepts, with the essence, or substantial form, as the unobserved, or latent, variable.

structures,¹⁷ because in these structures u is a confounding factor to determine the causal influence of k on e in terms of conditional probabilities:



To see potential empirical advantages of the causal power account, we should have examples with the above causal structures that are predicted to be true or acceptable on the probabilistic measure $\Delta^*P_k^e$, but false on the causal measure p_{ke} . We are not sure whether such examples exist for standard BP generics, but we will discuss some cases, leaving it to the reader (and perhaps future empirical research) to have a more definite opinion.

A well-known example of structure (i) involves yellow fingers (k) and lung cancer (e). It used to be the case that cigarettes had filters that caused smokers to get yellow fingers. We know by now that smoking also causes lung cancer. It follows that many people that have yellow fingers get lung cancer, and thus that $\Delta^*P_k^e$ (and $P(e|k)$) is high. But, obviously, getting lung cancer is not due to having yellow fingers, i.e., in this causal structure $p_{ke} = 0$. It is smoking (u) that causes both. The question is whether the following generic is nevertheless true:

(15) People with yellow fingers develop lung cancer.

If the answer to this question is negative, (15) is false, although $\Delta^*P_k^e$ and $P(e|k)$ are high. If so, this example shows an empirical advantage of the causal account.

Examples of type (i) are easy to analyze (though see section 5), because there is no causal relation at all between k and e . Causal structures (ii) and (iii) are different in this respect. Still, it will be the case in these structures that p_{ke} is low while $\Delta^*P_k^e$ is high, in case p_{ke} is low and $P(u|k)$ and p_{ue} are high. There are many examples like that. Suppose, for instance, that women drink significantly more tea on a regular basis than men and that it is somewhat better to drink tea than to drink, say, coffee. In many countries it is also the case that women have a higher life expectancy than the average life expectancy. Thus, there will be a positive correlation between ‘drinking tea’ and ‘higher than average life expectancy’. We wonder, however, whether this by itself makes the following generic true.

(16) People that drink tea regularly have a higher than average life expectancy.

If this generic is taken to be false, we take it that this is due to the fact that the direct causal effect of drinking tea on life expectancy is just small.

In section 4.2.1 we saw that in causal structures of the form (iii), there exists a difference between the *direct* power of k to cause e , $p_{ke}^{direct} = P(e_{-u}|k)$, and the *total* power to do so, $p_{ke}^{total} = p_{ku} \times p_{ue} + p_{ke}^{direct}$. It is natural to define for such cases a third measure of causal

¹⁷To be sure, there are many more complicated causal structures with various variables where this will be the case. To focus discussion, however, we look only at these simple cases.

power as well: *indirect* causal power: $p_{ke}^{indirect} = p_{ku} \times p_{ue}$. We hypothesize that these different measures are relevant for the analysis of generic sentences. More in particular, we think it is natural that we take generics of the form ‘*k* are *e*’ to be true if and only if p_{ke}^{direct} is high, at least if it is natural that both $p_{ke}^{direct} > 0$ and $p_{ke}^{indirect} > 0$, as in figure (iii).¹⁸ Consider generic sentences like:

- (17) a. Jews are capitalists.
b. Jews are communists.

It is true (or let us assume so) that relatively many Jews in Germany in the first half of the 20th century were working in the financial sector (and thus are capitalists). It is also true (or let us assume so) that relatively many Jews were (famous) communists during this period. So, according to the associative analysis discussed in section 2, these generic sentences are predicted to be true in that period (under our assumptions). Although the Nazi regime used both types of generics in their propaganda against the Jews (cf. O’Shaughnessy, 2016), it is now almost universally agreed that both sentences are (and were) false. The reason is, or so it seems, that the fact that relatively many Jews were working in the financial sector, or became communists, was *not directly* due to these people being Jewish. Instead, the reason was *indirect*: relatively many Jews were working in the financial sector (in Europe), for instance, because in contrast to the Jews, the catholic majorities in these countries were not allowed to loan money for a profit, although the need to borrow money was there. Moreover, the Jews were forbidden to occupy many other professions. Similarly, relatively many Jews were communists, because (perhaps because of the first reason) relatively many Jews were financially well to do, and could afford their children to receive a good education. It was this high education that made many young students feel that there was something wrong in European capitalists societies (or so one could argue). Suppose, if only for the sake of argument, that $\Delta^*P_k^e$ is high, with ‘*k*’ denoting Jews and ‘*e*’ denoting being a capitalist (or a communist). Suppose also that this is due mainly to $p_{ke}^{indirect}$ being high, because p_{ke}^{direct} is low. We might think of p_{ke}^{direct} to measure the amount of Jews that *would* have become capitalists even if they would have had the same possibilities as non-Jewish Europeans. Stating it somewhat differently, p_{ke}^{direct} measures the amount of Jews that became capitalists because of their individual intention to become one, while $p_{ke}^{indirect}$ measures the amount of Jews that became capitalists because there was (almost) nothing else to choose. If such are (or were) the facts, and we generally take (17a) to be false, it shows (or would show) another empirical advantage of the causal account.¹⁹

5. The asymmetry of causal powers

Let us now look at one of the most obvious predicted differences between the associative analysis based on $\Delta^*P_k^e$ and the causal analysis based on p_{ke} : causality is crucially *asymmetric* (or so we assume), while correlations need not be. This is similar to causal versus non-causal analyses of counterfactuals. Whereas Lewis’s (1973) analysis of counterfactuals is not necessarily asymmetric, more recent causal analyses that follow Pearl (2000) are. As a result, these causal

¹⁸This latter condition is meant to rule that we can’t say anymore that smoking causes cancer, because this line of causation is necessarily mediated by the generation of tar in lungs due to smoking.

¹⁹We take it that the following type of examples could be given a similar analysis:

- (i) a. Blacks are good in athletics.
b. Moroccans are good at soccer. and
c. Young male Moroccans are criminals.

analyses have to explain how to account for so-called ‘backtracking counterfactuals’ like ‘If she came out laughing, her interview went well’, counterfactuals in which the consequent cannot have been caused by the antecedent because the latter came later in time than the former.

Suppose we have a causal structure of the form $k \rightarrow e \leftarrow u$. It is well possible that in such cases $\Delta^*P_e^k = \frac{P(k|e) - P(k|\neg e)}{1 - P(k|\neg e)}$ has a high value, meaning that generics of the form ‘Objects of type e are (generally) of type k ’ are true in such circumstances according to the non-causal analysis discussed in section 2. On the causal analysis presented above, however, for the same sentence to be true, it has to be the case that the causal power of e to produce k , p_{ek} , has to be high. That, however, is impossible. To see this, recall that we defined p_{ke} in section 3 as $P(e|k, \neg u)$, where u denotes the disjunction of all potential causes of e different from k . To define p_{ek} this means that we should now consider $P(k|e, \neg b)$, where b denotes the disjunction of all alternative causes of k different from e . Assuming that k does not occur without a cause and that e is not a cause of k , it is obvious that $p_{ek} = P(k|e, \neg b) = 0$.²⁰ Thus, an important empirical consequence of our investigated analysis is that generic sentences are predicted to be **asymmetric**.

At first, this seems like an obviously false prediction: both of the following two generics seem true:

- (18) a. People that are nervous smoke.
b. People that smoke are nervous.

Perhaps such examples simply show that causality is not *semantically* relevant for the analysis of generics, it is at most relevant for *pragmatics*: people take, perhaps wrongly, generics to say something about causal powers. Perhaps. But even then we would need a causal analysis for (18b) within pragmatics. We believe that we can provide a causal analysis for generics like (18b). But there is a price to be paid: **ambiguity**. Although most generics of the form ‘ ks are e ’ are true because of the causal power of ks to produce e , others are true because of the causal power of e -ness to produce k . To give a more detailed account of the second reading of generics, we will define the probability that, given e , e is due to k , $P(k \rightsquigarrow e|e)$. Given that we derived before that in our causal structure $k \rightarrow e \leftarrow u$, objects of type e are caused by k with probability $P(k) \times p_{ke}$, the probability that, given e , e is due to k is

$$(19) \quad P(k \rightsquigarrow e|e) = \frac{P(k) \times p_{ke}}{P(e)}.^{21}$$

Notice that in causal structure $k \rightarrow e \leftarrow u$ this value can be positive and high, while $p_{ek} = 0$. Thus, on our ambiguity analysis we will say that although most generics of the form ‘Objects of type k are (generally) of type e ’ are true because p_{ke} is high, others are true because $P(e \rightsquigarrow k|k)$ is high. Observe that in contrast to p_{ke} , the value of $P(e \rightsquigarrow k|k)$ depends crucially on the base rates of k and e , making the latter less ‘stable’ than the former.²²

We have argued above against analyses of generics that required generics to be ambiguous. But

²⁰Alternatively, we might follow Pearl (2000) and measure the relevant causal power in terms of intervention as follows $P(k_e|\neg k, \neg e)$. But because in this type of causal structure intervention of e doesn’t influence the probability of k , and because now $\neg k$ is taken to be true, this means that $p_{ek} = P(k_e|\neg k, \neg e) = 0$, just as it should be.

²¹See Cheng et al. (2000). Notice that in case k is the only (potential) cause of e , $p_{ke} = P(e|k)$. In that case it immediately follows that $P(k \rightsquigarrow e|e) = \frac{P(k) \times P(e|k)}{P(e)} = \frac{P(k \wedge e)}{P(e)} = P(k|e)$.

²²Perhaps this explains why generics expressed in the ‘causal order’ are more natural than the others.

does the current ambiguity analysis not leave us with the same problem? Not really: in case one takes over Cheng's independence assumptions by means of which she can estimate the causal power, one can show not only that $p_{ke} = \Delta^*P_k^e$, but also that $P(e \rightsquigarrow k|k) = \Delta^*P_k^e$. Thus, as far as estimation is concerned, there is no ambiguity at all.

In this section we saw that a causal power analysis of generics has consequences for the analysis of generics in general. Whether an ambiguous, but perhaps more insightful, analysis of generics in terms of causal powers is to be preferred to a uniform, though less explanatory, analysis in terms of (stable) correlations, we must leave to the reader.

6. Conclusion and Outlook

The goal of this paper was to see to what extent a causal power analysis of generics is defensible. We have seen that such an analysis is quite appealing in the following sense: it **explains** why under natural circumstances a generic of the form 'ks are e' is true iff the measure $\Delta^*P_k^e$ is high, an analyses that was proposed before for empirical reasons. This explanation also has the conceptually appealing feature that it seems to align with our actual thinking. It forces us to look for suitable alternative potential causes and the relevant causal structures in which they are engaged. For instance, if two kinds both exhibit the same properties, it tries to come up with a common cause explanation. This forces one to look for 'deeper' analyses than a regularity analysis does. We feel, with Cartwright (1989), that this is also the way science works. Moreover, the causal analysis also gives rise to different empirical predictions in other than the 'natural' circumstances: (i) under various conditions generics of the form 'ks are e' are seen to be true, or acceptable, although $\Delta^*P_k^e$ is low, and (ii) it explains why some other generics are intuitively false, although $\Delta^*P_k^e$ and $P(e|k)$ are high. Moreover, we have seen that in various circumstances high causal power comes down to high (stable) conditional probability, which according to many authors (e.g. Cohen, 1999) is the reason why most generics are true.²³

It has to be admitted, though, that the empirical predictions of the causal power analysis of generics gives are not obviously better than those of the correlation-based analysis,²⁴ and that the analysis gives rise to some new problems as well. It forces us, for instance, to assume that generics – like counterfactuals – are ambiguous (although we have argued that this is less of a problem than it might seem at first). Another potential disadvantage of the causal power approach is that its resulting empirical predictions depend very much on what we take the alternative causes and the causal structure to be. Although the causal power theory might be seen as too 'subjective' as a result of this, we think that this is actually a fact of life: the discussions about whether sentences like 'Blacks are good in athletics' are true are very much about what the causal structure is taken to be.

In this paper we have investigated whether a *semantic* analysis based on causal powers is defensible. What the above doubts about this analysis shows, one might argue, is that a causal view should play a role only in *pragmatics*. In fact, this seems to be the view defended by

²³Another pleasing consequence is, that just like episodic sentences, generic sentences are on a causal power analysis predicted to be true just because a certain fact obtains. Because we assumed generics to have truth-makers (the causal powers) that are independent of the base rates, we predict—in contrast to purely probabilistic analyses—that generics express propositions and can be used in embedded contexts, like in 'Countries that do not honor women's rights, do not honor general human rights'.

²⁴Note that neither of these analyses seem suited all by itself to account for normative generics like 'Boys don't cry'.

Haslanger (1987) and Leslie (2013): the generic ‘Women are submissive’ should be avoided not so much because it is not true, but rather because it gives rise to the false suggestion that the generic is true for the wrong causal reasons, i.e., because of what it is to be a woman. In our words, ‘Being a woman’ is taken to be the *direct cause* of behaving submissive. One way to implement this suggestion is to claim that generics have truth conditions based on correlations, but that many people assume that these correlations are the way they are because of their wrong *essentialist*’ reading of generics. We have suggested in section 4.1 that if essences play a key role in the causal interpretation of generics, causal power reduces naturally to conditional probability. Although this might lead to a somewhat stronger reading of generics than the one using Δ^*P , it doesn’t lead to the much stronger interpretation that Haslanger and Leslie object to. Many proponents of a causal power view of regularities (Harre and Madden, 1975; Ellis, 1999), however, have something stronger in mind: the regularities are not just causal, but are taken to be (metaphysically) *necessary* (whatever that might mean exactly). It is exactly against this latter strong — and we think wrong — essentialist’ view of generics that Haslanger (1987) and others warn us. Haslanger argues—just like Barth (1971) before her—that because generic sentences like ‘Women are submissive’ and ‘Bantu’s are lazy’ are taken to say something about the essence of, or of the real, women and Bantus, they have their malicious social impact: they introduce prejudices to children, strengthen existing ones, and are excellent strategic tools for propagandists because they are immune to counterexample: any non-submissive woman is not a real woman. We think, however, that once the connection between causal powers (or essentialism) and necessity is given up, some of Haslanger’s complaints against the use of generics lose their force. It still leaves open, however, the idea that causal powers should be used in pragmatics, to account for the appropriateness of generic sentences, rather than in semantics, to account for their truth.

References

- Barth, E. (1971). *De Logica van de Lidwoorden in de Traditionele Filosofie*. Leiden: Universitaire Pers Leiden.
- Bird, A. (2007). *Nature’s Metaphysics. Laws and Properties*. Oxford: Oxford University Press.
- Cartwright, N. (1989). *Nature’s Capacities and their Measurement*. Oxford: Oxford University Press.
- Cheng, P. (1997). From covariation to causation: A causal power theory. *Psychological Review* 104, 367–405.
- Cheng, P., L. Novick, M. Liljeholm, and C. Ford (2000). Explaining four psychological asymmetries in causal reasoning: Implications of causal assumptions for coherence. In J. Cambell (Ed.), *Topics in Contemporary Philosophy, Volume 4: Causation and Explanation*, pp. 1–31. Cambridge, MA: MIT Press.
- Cohen, A. (1999). *Think Generic! The Meaning and Use of Generic Sentences*. Stanford: CSLI Publications.
- Danks, D. (2014). *Unifying the Mind. Cognitive Representations of Graphical Models*. Cambridge: Cambridge University Press.
- Ellis, B. (1999). Causal powers and laws of nature. In H. Sankey (Ed.), *Causation and Laws of Nature*, pp. 19–34. Dordrecht: Dordrecht.
- Gelman, S. (2003). *The Essential Child*. New York: MIT Press.
- Glymour, C. (2001). *The Mind’s Arrows*. Cambridge, MA.: MIT Press.

- Good, I. (1961). A causal calculus. *British Journal for the Philosophy of Science* 11, 305–318.
- Greenberg, Y. (2003). *Manifestations of Genericity*. London: Routledge.
- Harre, R. and E. Madden (1975). *Causal Powers: A theory of natural necessity*. Oxford: Basic Blackwell.
- Haslanger, S. (1987). Ideology, generics and common ground. In C. Witt (Ed.), *Feminist Metaphysics: Essays on the Ontology of Sex, Gender and the Self*, pp. 179–207. Dordrecht: Springer.
- Keil, F. (1989). *Concepts, kinds and cognitive development*. Cambridge, MA: MIT Press.
- Lawler, J. (1973). *Studies in English Generics*. Ann Arbor: University of Michigan Press.
- Leslie, S. (2008). Generics: cognition and acquisition. *The Philosophical Review* 117, 1–47.
- Leslie, S. (2013). Essence and natural kinds: when science meets preschooler intuition. In T. Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology*, 4, pp. 108–166. Oxford: Oxford University Press.
- Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.
- O’Shaughnessy, N. (2016). *Selling Hitler: Propaganda and the Nazi Brand*. Oxford: Oxford University Press.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. San Mateo, California: Morgan Kaufman Publishers.
- Pearl, J. (2000). *Causality: Models, Reasoning and Inference*. Cambridge: Cambridge University Press.
- Pearl, J. (2014). Interpretation and identification of causal mediation. *Psychological Methods* 19, 459–481.
- Preacher, K. and K. Kelley (2011). Effect sizes measures for mediation models: Quantitative strategies for communicating indirect effects. *Psychological Methods* 16, 93–115.
- Schanks, D. (1995). *The Psychology of Associative Learning*. Cambridge: Cambridge University Press.
- Shep, M. (1958). Shall we count the living or the dead? *New England Journal of Medicine* 259, 1210–1214.
- Shoemaker, S. (1980). Causality and properties. In V. Inwagen (Ed.), *Time and Cause*, pp. 109–135. Dordrecht: Reidel.
- Skyrms, B. (1980). *Causal Necessity*. New Haven: Yale University Press.
- Slooman, S. (2005). *Causal Models. How People Think about the World and Its Alternatives*. Oxford: Oxford University Press.
- Spirtes, P., C. Glymour, and R. Scheines (2000). *Causation, Prediction and Search*. Cambridge MA: MIT Press.
- Strawson, G. (1989). *The Secret Connection: Causation, Realism and David Hume*. Oxford: Oxford University Press.
- Tversky, A. and D. Kahneman (1980). Causal schemas in judgments under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology*, pp. 49–72. Erlbaum: Hillsdale.
- van Rooij, R. and K. Schulz (to appear). Generics and typicality: a bounded rationality approach. *Linguistics and Philosophy*, to appear.