A causal semantics of IS generics

van Rooij, R.; Schulz, K.

DOI
10.1093/jos/ffz023

Publication date
2020

Document Version
Final published version

Published in
Journal of Semantics

License
Article 25fa Dutch Copyright Act

Link to publication

Citation for published version (APA):
van Rooij, R., & Schulz, K. (2020). A causal semantics of IS generics. *Journal of Semantics*, *37*(2), 269-295. https://doi.org/10.1093/jos/ffz023

Journal of Semantics, 37, 2020, 269–295
doi:10.1093/jos/ffz023
Advance Access Publication Date: 24 February 2020
Articles

OXFORD

A Causal Semantics of IS Generics

Robert van Rooij
Institute for Logic, Language and Computation (ILLC), University of Amsterdam

Katrin Schulz
Institute for Logic, Language and Computation (ILLC), University of Amsterdam

First version received 11 July 2018; Second version received 2 December 2019; Accepted 5 December 2019

Abstract

The felicity, or acceptability, of IS generics, i.e. generic sentences with indefinite singulars, is considerably more restricted compared to BP generics, generics with bare plurals. The goal of this paper is to account for the limited felicity of IS generics compared to BP generics, on the one hand, while preserving the close similarity between the two types of generics, on the other. We do so by proposing a causal analysis of IS generics, and show that this corresponds closely with a probabilistic analysis of BP generics.

1 INTRODUCTION

Generics are sentences that express basic generalities without the use of an explicit quantifier. Generic sentences come in various sorts: they can be expressed using a bare plural (BP), a singular indefinite (IS), and a definite description. In this paper we will be concerned with the meaning of generics in general, and in particular with the relation between the meaning of BP generics like (1-a) and IS generics like (1-b).

(1) a. Tigers have stripes.
b. A tiger has stripes.

Although (1-a) and (1-b) seem to express the same content, Lawler (1973), Burton-Roberts (1976, 1977), Krifka (1987, 2012), Cohen (1999, 2001), Greenberg (2003, 2007), Leslie (2009), and others have argued that there still exists a crucial distinction between the two types of generics that should be accounted for. In this paper we will argue that BP and IS generics are very similar in that they both mainly express inductive generalizations. With Lawler, Greenberg and Leslie et al. we will argue that whereas BP generics could already be true due to a merely statistical correlation between the denotations of subject and verb

The ideas behind this paper was presented at several workshops and conferences. We thank the audiences for their comments and suggestions. We also would like to thank the reviewers of this paper for their comments and challenging questions. We would also like to thank the associate editor.

phrase, IS generics *demand* a 'principled connection'. The main innovation of this paper is to argue that this 'principled connection' is mostly a causal one and to propose a particular implementation of this view.

This paper is structured as follows: in the following section we will briefly motivate a recently proposed descriptive probabilistic analysis of BP generic sentences. In section 3 we will look at what is particular to IS generics, and at previous proposals of how to account for this. In section 4 we will provide a background in causal models, and show how generics —and in particular IS ones —can be given a causal analysis. In section 5 we will discuss some generalizations of this causal analysis. Section 6 concludes the paper.

2 A PROBABILISTIC-BASED ANALYSIS OF GENERICS

Bare Plural (or BP) generic sentences like 'Birds fly' and 'Tigers are striped' (which we take to have the form '*x*s are *y*') are sentences that, by their very nature, express useful generalizations. The main question addressed in the literature is about the type of generalization. First, generic sentences are clearly not universally quantified sentences: although not all birds fly (penguins don't), (2) is a good generic sentence that most people consider true, or acceptable.[1]

(2) Birds fly.

Indeed, this is one of the most typical features of generic sentences: they express generalizations that allow for exceptions. But it also need not be the case that almost all, or *most* *x*s have to have feature *y* in order for the generic '*x*s are *y*' to be true or acceptable:

(3) a. Birds lay eggs.
    b. Goats produce milk.

Although people normally accept (3-a) and (3-b), it is not the case that the majority of birds or goats have the relevant feature; only the adult female birds and goats do! Moreover, even if most *x*s are (or are taken to be) *y*, the corresponding generic sentence still doesn't have to be true or acceptable, as exemplified by sentences such as the following (from Carlson, 1977):

(4) a. ∗ Books are paperbacks.
    b. ∗ People are over three years old.
    c. ∗ Bees are sexually sterile.

According to a natural alternative quantificational proposal, a generic is true/acceptable exactly if all, or most, *normal*, or *relevant* *x*s are *y*, or perhaps if all x, if they were normal x, would also have property *y* (cf. Asher & Morreau (1995)). There are at least two problems with such an analysis. First, the above problem(s) for the majority account would have to be tackled.[2] Second, and more importantly, without an independent analysis of what it is

1 In this paper we will leave it open whether generics have truth-conditions, or merely have acceptability conditions.

2 See Nickel (2009) for an (unsuccessful, we believe) attempt. Nickel concentrates attention to another problem of such an analysis: the fact that (i) is a good generic:

(i) Elephants live in Africa and Asia.

to be a normal or relevant x, such an analysis hardly makes any empirical prediction (cf. Krifka *et al.* (1995)).

Following Cohen (1999, 2004), we believe that there is a lot more to say for a simple and intuitive (high) majority-account than is standardly acknowledged. First, there exists a natural way to account for (3-a) and (3-b) in these terms after all. As argued for by Cohen (1999), to interpret a generic, we should take into account the alternatives of the predicate term. For (3-a), for instance, we should take into account $Alt(lay\ eggs)$, which naturally is $\{lay\ eggs, give\ live\ birth\}$. Cohen (1999, 2004) accounts for generics making use of an 'objective' probability function. (2), for instance, is taken to be true if and only if the conditional probability $P(Fly|Birds)$ is higher than $\frac{1}{2}$. But for the analysis of (3-a), however, the probability function should now not range over *all* objects, but only over objects that satisfy at least one of the properties in $Alt(lay\ eggs)$. Because $\bigcup Alt(lay\ eggs) \approx Females$, a majority analysis could, or would, predict that (3-a) is true, or acceptable, just in case the majority of *female* birds lay eggs.[3] A similar move could be made with (3-b).

Second, to account for the intuition that (4-a)–(4-c) are bad generics, we can make use of Cohen's *homogeneity condition*. Rather than just demanding that $P(y|x)$, the conditional probability of y given x, is high, Cohen (1999, 2004) demands that the conditional probability of y given a set of xs should be high for each cell of a salient partition $\{x_1, ..., x_n\}$ of x. Thus, each of $P(y|x_i) \cdots P(y|x_n)$ should be high. Concentrating on (4-c), for instance, a salient partition of bees into queens (female), workers (female) and drones (male) will correctly predict that (4-c) is bad, because neither queens nor drones tend to be sterile. Cohen provides a similar explanation for the other examples as well. That seems very promising.

Still, Leslie (2008) has persuasively argued that the condition of homogeneity not only explains away bad generics, but good ones as well. Why, for instance, is 'Bees reproduce' true on Cohen's salient partition of bees? The same holds, even, for the standard example 2 'Birds fly'. Why shouldn't this generic not predicted to be simply false, or unacceptable, by a bi-partitioning of birds into Penguins, on the one hand, and all the other types of birds, on the other? Clearly, for the homogeneity condition to work, there have to be constraints on what good partitions are.

We do think that a homogeneity condition is at work for the analysis of generics. Cohen (2004) provides a partial motivation for his condition in terms of methods of categorization. By taking generics to express *inductive generalizations*, there is also another motivation. The reason is that for inductive generalizations to be reliable, the generalizations have to be *stable*.

> No one supposes that a good induction can be arrived at by merely counting cases. The business of strengthening the argument chiefly consists of determining when the alleged association is *stable*, when the accompanying conditions are varied. (Keynes (1952), p. 393)

This is a problem for a standard normality account, because no elephant lives in both continents. Nickel proposes that generics like 'xs are y' are true if there is a subkind x' of x for which the generic 'x's are y' is true on its normality reading. With Hoeltje (2017), however, we think this proposal is too weak. A much more natural analysis for (i), we believe, would be to claim that the sentence should be given a cumulative reading, recently proposed by James Kirkpatrick in still unpublished work.

3 Although such an analysis seems natural, it is not made by Cohen (1999).

Something like the homogeneity condition is used exactly to account for this intuition. Skyrms (1980) proposes to account for stability in terms of *invariance* under conditionalization. Take S to be a set of contextually given sentences logically independent of y, thus consistent with y and $\neg y$. Then $P(y)$ is (relatively) *invariant to* S iff $\forall s \in S : P(y|s) \approx P(y)$. This notion of invariance can be generalized to conditional probabilities, $P(y|x)$, as well, but then we have to limit ourselves to members of S that are consistent with $x \wedge y$ and $x \wedge \neg y$. Skyrms refers to stable, or invariant, probabilities as *propensities*.

Demanding high propensity, i.e., high *stable* probability, instead of just high probability, we can immediately explain why the partition {*Penguins*, *other birds*} is not a good partition to base the generalization 'Birds fly' on. It simply doesn't satisfy the conditions for being an appropriate set with respect to which the conditional probability $P(Fly|Bird)$ has to be stable, because one cell of the partition is inconsistent with 'Fly'. On the other hand, Skyrms' condition of approximate invariance still explains why the generics (4-a)–(4-c) are bad, just like Cohen's (1999; 2004) original condition of homogeneity did.

Still, Leslie (2008) objects: how can a generic like 'Bees reproduce' still be true on such an analysis? Skyrms' condition doesn't rule out that the relevant conditional probability should be high for all members of the partition {*male bees*, *female bees*}, and most, though not all, female bees (which constitute the vast majority of bees) are sterile (or at least sexually inactive). We agree, but we also believe that there is another reason why this generic is true, or good, nevertheless. It is well-known that plurals and mass nouns have, besides a 'distributive', also a '*collective*', or 'semi collective', interpretation. On such a (semi-) collective interpretation of sentence '(the) xs y', it doesn't have to be the case that all (minimal) parts of x have to have feature y for the sentence to be true, it is only required that the whole group — or better, larger subsets of this group — does so. Notice that bees, like ants, are very small insects, and that most of us, or that we most of the time, do not individuate such insects. In fact, we take it, when we think about reproduction of bees and ants, we think of these small insects most of the time as *collections*, e.g. as swarms, nests, hives, or armies. Moreover, although in English, 'bee' and 'ant' are count nouns, their counterparts in languages such as Russian, Welsh, and Dagaare are actually mass-nouns (cf. Grimm (2009)). This suggests that it is at least natural to view bees and ants primarily as collections. As observed already by Krifka *et al.* (1995), Cohen (1999) and others, some generic sentences (seem to) have a collective interpretation. But this solves our puzzle: 'Bees reproduce' is true, or acceptable, because it receives a (semi) collective interpretation, and as hives, bees generally *do* reproduce. (Notice that also on this semi-collective interpretation, the sentence is still interpreted as a generic.)

Still, a (stable) majority analysis of generics seems certainly problematic when accounting for the following examples.

(5) a. Ticks carry Lyme disease.
    b. Sharks attack bathers.
    c. Pit bulls are dangerous.

Intuitively, the above sentences are appropriate, although only 10% of ticks carry Lyme disease, only very few sharks ever attack bathers, and pit bulls normally don't do any harm. There are several strategies to deal with such sentences. According to one of them, what counts is not whether, for (5-a) for instance, the majority of ticks *actually* carry Lyme disease, but whether they *can* carry Lyme disease. But, obviously, this strategy overgenerates enormously: why is an example like 'Chairs are painted white' not true, just because chairs *can* be painted white? According to another strategy (the 'error strategy'), one might say that these sentences are actually false. But why, then, do so many people take them to be true? A major worry here is to determining what the data are: if (5-a)–(5-c) are generally taken to be true, what is it that makes the claim 'correct' that these sentences are *in fact* false? It cannot be that this is so because it is predicted by the theory, because the theory itself is based on intuitions of the language users.

Cohen (1999) proposed that generics like (5-a)–(5-c) should be interpreted in a *relative* way: (5-a) is true iff compared to other types of animals *relatively many* ticks cary the Lyme disease. Similarly for (5-b) and (5-c). In probabilistic terms this means that $P(C|T) > P(C)$ — or better $P(C|T \cap \bigcup Alt(C)) > P(C|is \bigcup Alt(C))$ — should hold with 'T' denoting the Ticks and 'C' standing for 'carry Lyme disease'.

Although appealing, there are several reasons to be dissatisfied with Cohen's proposal. One of the reasons why a (stable) majority account of generics is so natural, is that it motivates the inference of *instantiation*: if the propensity of xs having feature y is high, it is natural to assume, or act as if, an arbitrary x also has feature y. However, this type of inference hardly makes any sense when just relatively many xs have feature y. This, we take to be the *conceptual problem* of Cohen's relative reading. There is a, perhaps even more serious, *empirical problem* as well: as noted by Leslie (2008), how to prevent that also a generic like 'Dogs have three legs' is predicted to be true as well, given that dogs are pets, which, in contrast to other animals, do not die, or are not killed, if they loose a leg? The conclusion seems obvious: the conditions for being true on its relative reading just seem to be *too weak*.

Leslie (2008) proposed that (5-a)–(5-c) are (taken to be?) true because the predicates involved in these sentences are *horrific*, or *striking*. Although we don't agree with her analysis (which would make generics highly ambiguous), we think that an appeal to how 'striking' the predicate is to account for the truth, or acceptability, of these examples is correct.

One way to make use of this intuition is to say that (5-a) is a good generic, not because most ticks carry Lyme disease, but because learning that a significant portion of ticks carry Lyme disease is useful, or important. Learning this information is enough to take action: the *expected value* of taking action is higher than doing nothing, because the Lyme disease is dangerous. The same holds for (5-b) and (5-c). We could define the expected value of a generic of the form 'xs are y' as $EV(y|x) = P(y|x) \times V(x \wedge y)$, or perhaps $P(y|x) \times V(y)$, where $V(y)$ measures the value, impact, or strikingness of feature y.[4] Then we could say that the generic is true, or acceptable, just in case $EV(y|x)$ is high.

But what does it mean for $EV(y|x)$ to be high, given that $EV(y|x)$ has no clear minimum or maximum? Perhaps $EV(y|x)$ should be high compared to some natural alternatives of y, i.e., compared to $EV(\bigcup Alt(y)|x)$, or of alternatives of x, i.e., compared to $EV(y|\bigcup Alt(x))$. The first seems unnatural: using it would mean that any generic with 'horrific' features would almost automatically be true or acceptable. But that, again, is a clear false prediction. That leaves us to compare $EV(y|x)$ with $EV(y|\bigcup Alt(x))$, which we will abbreviate from now on by $EV(y|\neg x)$.

4 We assume that features with 'negative' connotation have more impact than those with positive ones. Still, also for features with negative connotation, $V(y) > 0$, and for such y typically $V(y) > 1$.

If for evaluating the truth, or acceptability, of 'xs are y' we should compare $EV(y|x)$ with $EV(y|\neg x)$, this naturally means that it should be the case that $EV(y|x)$ is significantly higher than $EV(y|\neg x)$, i.e., $EV(y|x) \gg EV(y|\neg x)$. Now suppose that $V(y) = 1$. In that case, the condition $EV(y|x) \gg EV(y|\neg x)$ comes down to $P(y|x) \gg P(y|\neg x)$, from which it follows that $P(y|x) > P(y|\neg x)$. It is easy to show that $P(y|x) > P(y|\neg x)$ iff $P(y|x) > P(y)$ which is not only a standard notion of statistical *relevance* (e.g. Adams (1998)), but also the condition Cohen (1999) required for a generic sentence to be true on its relative reading.[5] In case the value of y, $V(y)$, is high, i.e., in case y is a 'horrific' or 'striking' feature, it could already be the case that $EV(y|x) \gg EV(y|\neg x)$ if $P(y|x)$ is somewhat higher than $P(y|\neg x)$, because the difference is going to be multiplied by $V(y)$.[6] For a more 'neutral' feature, however, where $V(y) = 1$, $P(y|x)$ itself has to be significantly higher than $P(y|\neg x)$ for the generic to be true, or acceptable.

In van Rooij & Schulz (2019a) it is proposed that not just some, but *all* (BP) generic sentences should be analyzed by demanding that $EV(y|x) \gg EV(y|\neg x)$.[7],[8] This immediately explains why in contrast to (6-a), the generic (6-b) is acceptable, although only male lions have manes, and thus $P(male|lion) > P(have\ mane|lion)$.

(6) a. $*$ Lions are male.
    b. $\surd$ Lions have manes.

The claim that *all* generics should be based on the strengthened relative reading of Cohen (1999) is based not just on empirical grounds, but also on a conceptual one: Generic sentences are not only the simplest type of sentences that express generalities, they are also among the first type of sentences we **learn** to understand and use in childhood (see e.g. Gelman, 2003). This suggests that the meaning of generic sentences should be based on simple principles of *inductive learning* (cf. Leslie (2008)). According to a standard psychological analysis of inductive learning (in adults, children *and* animals), we learn to associate objects or events of type x with objects or events of type y iff $\Delta P_x^y =_{df} P(y|x) - P(y|\neg x) > 0$ (cf. Shanks (1995)), or better, when this difference is *stable* under various conditions (cf. Cheng & Holyoak (1995)). Obviously, ΔP_x^y (known in psychology as the measure of *contingency*) is positive just in case $P(y|x) > P(y|\neg x)$ iff $P(y|x) > P(y)$, which means that $\Delta P_x^y > 0$ is positive just in case Cohen (1999) would judge the generic of the form 'Objects of type x are of type y' true on its relative reading.

5 At least, if we ignore the alternatives.

6 Suppose it is only the value of the feature, $V(y)$, that counts. Then the difference is going to be multiplied by $V(y)$, because $EV(y|x) - EV(y|\neg x) = P(y|x) \times V(y) - P(y|\neg x) \times V(y) = [P(y|x) - P(y|\neg x)] \times V(y)$.

7 In the same paper it is suggested how also generics with an 'existential' reading (like in 'Indians $[\text{do}]_F$ eat beef' as denial to an earlier statement that Indians don't eat beaf) can be accounted for on this analysis.

8 For a closely related recent proposal, see Tessler & Goodman (2019). One important difference between the proposals is that for us utility, or value, is important, while for them it is not. They also don't make use of causality, which will be crucial in the main part of this paper. See van Rooij & Schulz (2019a) for a somewhat more detailed discussion of the differences.

It is noteworthy that the measure ΔP^y_x can account for the intuition that generics are true, or acceptable, because they express *stereotypical* properties of groups. Social psychologists (e.g. Schneider, 2004) have proposed two measures of stereo-typicality: $\frac{P(y|x)}{P(y|\neg x)}$ and $\frac{P(y|x)}{P(y)}$, and it is clear that both these measures are large just in case the difference between $P(y|x)$ and $P(y|\neg x)$ is high. Other psychologists (Tenenbaum & Griffiths (2001); Tenori *et al.* (2007)) have proposed these, or very similar, measures for modeling the *representativeness*, or *typicality* of y for x. Thus, by using ΔP^y_x we would in fact say that a generic of the form 'xs are y' is true or acceptable just in case y is a (stereo)-typical, or representative feature of xs.

We have argued above that for a generic of the form 'Objects of type x are y' to be true, or acceptable, the value of $P(y|x)$ should be *significantly higher* than $P(y|\neg x)$ (from now on in this paper we will ignore the 'striking' generics, and concentrate only on probabilities). But this by itself cannot account for the intuition that what counts as 'significantly higher' depends very much on the base rate of y, i.e., $P(y)$ (or, equivalently, $P(y|\neg x)$). In case $P(y)$ is low, i.e., when y is an uncommon feature, it seems that the difference between $P(y|x)$ and $P(y|\neg x)$ should be larger than when y is a common feature. This can be accounted for by demanding that for a generic of the form 'Objects of type x have feature of type y' to be true or acceptable, Shep (1958) measure of *relative difference*, $\Delta^* P^y_x$, has to be significantly higher than 0, or significantly higher than $\Delta^* P^z_x$, with z any (contextually) relevant alternative feature to y, and with

(7) $\Delta^* P^y_x =_{df} \frac{\Delta P^y_x}{1-P(y|\neg x)} = \frac{P(y|x)-P(y|\neg x)}{1-P(y|\neg x)}.$

$\Delta^* P^y_x$ is the ratio of the amount by which x increases the probability of y, ΔP^y_x, to the room available for increase, $1 - P(y|\neg x)$.[9]

Notice that in case $P(y|\neg x)$ is high, ΔP^y_x counts for much more than when $P(y|\neg x)$ is low. For example, if $P(y|\neg x) = 0.9$, $\Delta^* P^y_x = \mathbf{10} \times \Delta P^y_x$. One can also easily show that in contrast to ΔP^y_x, for the measure $\Delta^* P^y_x$ the value of $P(y|x)$ counts for more than the value of $P(y|\neg x)$, as it intuitively should.[10] For instance, it is easy to see that once $P(y|\neg x) < 1$, $\Delta^* P^y_x = 1$ just in case $P(y|x) = 1$, i.e., $\Delta^* P^y_x$ receives its maximal value just in case $P(y|x)$ does so.[11]

9 There exists another notion of 'relative difference', proposed by Crupi et al. (2007) to measure the value of confirmation of an hypothesis, that could be used as well: $\frac{P(y|x)-P(y)}{1-P(y)}$. At this point we don't have good arguments why we should prefer $\Delta^* P^y_x$ over this measure.

10 It is perhaps useful to note that $\Delta^* P^y_x$ is equivalent with $\frac{P(y|x)-P(y)}{P(\neg y \wedge \neg x)}$.

11 The demand that $EV(y|x) \gg EV(y|\neg x)$, or that $\Delta^* P^y_x \gg 0$ might still seem too weak. In van Rooij & Schulz (2019a) it is argued that the analysis is strengthened pragmatically, such that it *feels like* $P(y|x)$ is high. The strengthening has three sources: (i) in psychology it is known that 'striking' events are better remembered, and taken to have a higher probability than they actually have. Thus, if $EV(y|x)$ is high due to high $V(y)$ this gives rise to the 'feeling' that $P(y|x)$ is high; (ii) it is relatively well established (cf. Tversky & Kahneman's 'conjunction problem') that people confuse inductive support (measured in terms of something like ΔP^y_x or $\Delta^* P^y_x$) for conditional probability, meaning that if the former is significantly above 0, this is taken to mean that conditional probability of y given x is high; and (iii) notions of inductive support are typically taken to implicate a causal relation, which in turn is giving rise to an essentialist' meaning, suggesting that $P(y|x)$ is high. One might think of the present paper as arguing that for IS generics, the latter 'psychological' idea is taken to be part of semantics.

3 INDEFINITE SINGULAR (IS) GENERICS

Generic sentences can be expressed not only using bare plurals, like '*Dogs* have four legs', but also with indefinite singular noun phrases, as in '*A dog* has four legs'. Given that 'Dogs have four legs' and 'A dog has four legs' are both appropriate, and seem to allow in similar ways for exceptions, they seem to express the same content. Still, the linguist John M. Lawler observed in Lawler (1973) that sometimes the two types of generics come apart: although (8-a) can be appropriate, or acceptable, (8-b) (read generically) cannot.

(8) a. √ Rooms are square.
    b. ∗ A room is square.

Similarly, Oosterhof (2008), p. 66) observes that (9-a) is ok, but that (9-b) is not.

(9) a. √ Musical comedies (plays) are popular.
    b. ∗ A musical comedy (play) is (a) popular (play).

Greenberg (2003, 2007) observes that the felicity of IS generics is considerably more restricted than that of BP generics in other ways as well. She notes that IS, but not BP generics are infelicitous, or unacceptable, with subjects denoting extremely unnatural properties:[12]

(10) a. √ Norwegian students with names ending with 's' wear thick green socks.
     b. ∗ A Norwegian student with a name ending with 's' wears thick green socks. (odd as generic, fine as existential)

(11) a. √ Well-known forty-five-year-old teachers do not cook on Monday afternoons.
     b. ∗ A well-known forty-five-year-old teacher does not cook on Monday afternoons. (odd as generic, fine as existential)

She also notes that, IS, but not BP, generics are infelicitous with episodic predicates:

(12) a. ∗ An earthquake is/√Earthquakes are especially strong today.
     b. ∗ An Italian restaurant is/√ Italian restaurants are closed tonight.

To account for such facts, Burton-Roberts (1976, 1977) and Krifka (2012) argue that whereas BP generics just express empirical generalizations, IS generics express analytic (Burton-Roberts) or definitional (Krifka) truths, sentences that are true on the basis of linguistic rules or conventions.[13] Although Cohen (2001) agrees that IS generics can express linguistic conventions, he points out that IS generics are not limited to that: they can also express moral and legal rules, and rules concerning games as well.

(13) a. A gentleman opens doors for ladies.
     b. A good king is generous.
     c. A grapefruit costs 1 euro.
     d. A bishop moves diagonally (in chess).

12 One reviewer didn't see a strong contrast in the first example, though.

13 Burton-Roberts (1976) proposed that sentences like (i-a) have the same meaning as sentences like (i-b).

(i) a. A kangaroo is a marsupial.
    b. To be a kangaroo is to be a marsupial.

Cohen (1999, 2001) argues that in contrast to BP generics, IS generics can only express rules or regulations, which includes more than just linguistic rules.

Although we agree that there is an important difference between IS and BP generics, we believe that the proposals of Burton-Roberts, Krifka and Cohen of how to account for this are problematic. Consider first a generic like (14), which Krifka (2012) still treats as definitional.

(14) A donkey has 62 chromosomes.

Intuitively, donkeys are not *by definition* animals that have 62 chromosomes: men didn't first have the notion of 'having *x*-many chromosomes' and then defined what it is to be a donkey in terms of it. No, we always knew what a donkey is, but *discovered* that it has 62 chromosomes. Kripke (1980) and Putnam (1975) would say that (14) expresses a truth that is *necessary* but *a posteriori*. Thus, in distinction with analytic truths, the sentence does not express an *a priori* truth. Still, Krifka and Burton Roberts (might) argue, there is a sense in which (14) can be called 'definitional' after all. In terms of Leibniz's 1684 terminology, (14) does not express the *nominal* definition of donkeys, but it still gives the *real* definition of donkeys: (14) states what (part of) the real essence of donkeys is. We agree, but we also think that to treat both types of definitions simply in the same way (as suggested by Burton-Roberts and Krifka) would only obscure the insight of traditional philosophy that there is an important difference between the two types of definitions, an insight that was forcefully renewed by the arguments due to Kripke and Putnam: in contrast to the nominal definition 'A bachelor is an unmarried man', the real definition 'Water is H_2O', though also necessary, does not express an *a priori* truth.[14]

According to Kripke and Putnam, a sentence like (14) expresses an *a posteriori* truth, and thus an empirical generalization, a view that seems at odds with the view that IS generics can only express rules or regulations. Indeed, as admitted by Krifka (2012), the following IS generics seem to express descriptive generalizations as well:

(15) a. A trout can be caught by many different methods.
b. A hedgehog makes a good pet.
c. A poodle should be clipped by a professional groomer.
d. A madrigal sounds best with all the voice-parts doubled.
e. A refrigerator costs about 1000 euros.

Burton-Roberts (1976, 1977), Krifka (2012) and Cohen (1999, 2001) (but see also Krifka (1987)) assign to IS and BP generics two completely different semantic structures. Such theories will have problems capturing the fact that also IS generics can express descriptive, and counterfactual supporting, generalisations. Moreover, such theories will have trouble explaining why IS generics can allow for exceptions (e.g., 'A dog has four legs').

Our proposal of how to account for IS generics will be more in line with the views of Lawler (1973) and Greenberg (2003, 2007) and with the experimental data found by Leslie (2009). Lawler (1973) argues that the more restricted felicity, or acceptability, of IS generics shows that a generic expressed with the indefinite singular requires a *stronger*

14 Krifka (2012) follows Stalnaker (1978) and others showing that the framework of two-dimensional modal logic can be used to account for the necessity of definitional sentences. That might be true, but *a posteriori* necessity is still something quite different from *a priori* necessity.

relation between the noun phrase and the property that is attributed to it, than when the generic is expressed with a bare plural. He argues that this stronger relation should involve 'essences'. Greenberg (2003) agrees to a large extent with Lawler and argues that in contrast to a BP generic, the truth of an IS generic requires a 'principled connection' between the kind of object denoted by the noun phrase and the property attributed to it.[15] Notice that both views are compatible with the idea that IS generics can express empirical generalizations.

On the basis of experimental evidence, Prasada & Dillingham (2005) make a distinction between two types of ways properties can be 'true of' kinds. Either due to a 'principled connection' (as in 'Cheetahs run fast', or 'Dogs have four legs') or due to a more accidental relation (as in 'Dogs wear collars'). We agree with Carlson (2009) when he claims that Lawler and Greenberg, on the one hand, and Prasada & Dillingham, on the other, were really 'talking much about the same thing'. This doesn't mean, of course, that when a generic is phrased with the bare plural that such a 'principled connection' does not exist. Indeed, the experimental study of Leslie (2009) indicates that sentences that express 'principled' generic assertions are nearly equally well expressed with bare plurals as with indefinite singulars (Dogs have / A dog has four legs and Ducks lay / A duck lays eggs), whereas generic assertions that just express a statistical generalization (Barns are red vs. A barn is red) are better expressed with bare plurals.

These phenomena suggest that in contrast to BP generics, IS generics are necessarily about *kinds*. But, of course, it is not very clear what, exactly, kinds are. Moreover, given an example like (15-e), kinds are not necessarily *natural* kinds. But the suggestion is on the right track: kinds are *intensional* rather than just extensional objects: they can, somehow, be identified, independently of how the world actually is. This, in turn, indicates that IS generics cannot be analyzed in terms of, just, actually observed, or observable, frequencies. For the truth, or acceptability, of IS generics, there must exist a 'principled connection'.

What does it mean for there to be a 'principled connection'? Although we don't agree with Krifka's (2012) 'definitional' analysis of IS generics, we do believe that he points to a natural idea in the following quote:

> Descriptive generalizations are typically expressed by bare plurals because they typically rely on observing many instances; definitional statements are typically expressed with indefinite singulars because it can be determined with single individuals whether or not they fit to the defining properties. (Krifka (2012), p. 389)

This quote captures not only the natural intuition that the use of a bare plural indicates that a BP generic expresses a generalization the truth, or acceptability, of which can only be established by observing many instances. It also proposes that the use of the *indefinite* in an IS generic signals that the expressed generalization (about a kind) can be established already on the basis of the observation of a single individual (of the kind). We think that both ideas are natural, even if it is acknowledged that IS generics can also express *empirical* generalizations. At first, such a proposal concerning (certain) inductive generalizations will, no doubt, sound almost like a contradiction in terms. Still, it makes a lot of sense. Take any particular kind of metal or liquid. Once we observed that one piece of a particular metal

15 We agree, but we also agree with Krifka (2012) that she implements this intuition in a rather roundabout way. It should be noted that Greenberg (2007) argues that both IS and BP generics express what she calls 'in virtue of' generalizations. The difference between the two kinds of generics only indicates whether this 'in virtue of' property is specified, or left vague.

melts under normal circumstances at a particular temperature, we inductively conclude that this will be the case for *all* pieces of this particular metal (cf. Sankey (1997)). The same holds for boiling temperatures of liquids.[16] In a corpus study on Dutch and Belgian texts, Oosterhof (2008) shows that (definite singulars and) indefinite singulars are used more frequently with animal names than with nationality names, indicating that in these texts IS generics are less frequent with social kinds than with natural kinds. This is consistent with the hypothesis under consideration, given that the members of natural kinds are normally more alike, or homogenous, than members of social kinds. In fact, Hume's (in)famous problem of induction is taken to be less of a problem by scientists than it is for (many) philosophers, because in contrast to philosophers, scientists uncritically take objects (of the same natural kind) to be so similar to one another that induction is taken to be reliable (see, e.g., Sankey (1997)). In this respect — or so we think — natural language shows that common sense agrees more with the (perhaps naive) scientific metaphysical outlook, than with the Humean more sceptical philosophical one.

Unfortunately, the above hypothesis concening IS generics can't be quite right. First, although all female ducks have somewhat boring grey/brown feathers, the IS generic (16) still seems, to us, appropriate.

(16) A duck has beautiful feathers.

Perhaps less controversially and more importantly, the following IS generic is also ok, although it obviously allows for exceptions.[17]

(17) A hungry dog is dangerous.

Thus, it is not the case that IS generics about kinds can only express properties *all* the members of this kind have. This seems to contradict the above hypothesis that empirical generalizations about kinds can be expressed by IS generics only if the generalizations can be made already on the basis of the observation of a single member of the kind.

We propose that what is signaled with the use of the indefinite in an IS generic is not on the basis of what a generalization can be established, but rather whether or not knowledge about observed frequencies can be used reliably for individual cases. It is well-known that on a standard frequency-based analysis of probabilities, the observed probabilities cannot say anything of any real value concerning individual cases. If the relevant frequency of outcome 'heads' for tosses of a coin is x, we can't conclude from this that the chance of this outcome for each individual toss will be x as well, simply because relative frequencies do not exist for single tosses of a coin. It was exactly to make sense of probabilities of individual cases that Popper (1979) proposed his *propensity interpretation* of probabilities. In contrast to Skyrms (1980) notion of propensity, on Popper's interpretation, stable long-

16 Of course, the generalization *that* we can infer from a single instance to a whole group for metals and liquids concerning melting and boiling temperatures cannot be determined from single instances. Not all inductive generalizations work this way.

17 We don't take this generic to be ok due to the striking predicate 'dangerous'. In fact, the IS variant of a striking generic like 'Sharks attack bathers', i.e., 'A shark attacks bathers' is typically taken to sound awful. In an earlier version of this paper, we used 'A dog is dangerous, when it is hungry'. A reviewer complained, claiming that it is well known that odd IS generics are often perfectly fine when modified by a when-clause. That might be so, but we don't see the difference between the when-sentence and (17).

run frequencies are a *manifestation of* something deeper: invariant *single-case* probabilities. Thus, propensities are not the relative-frequencies themselves, but rather the *causes of* the observed relative frequencies.[18] In this paper we will make use of Pearl's (2000) causal framework. We propose that the relevant 'principled connections' that we were looking for are (mostly) *causal* connections. Making use of Pearl's (2000) causal framework allows us to build a bridge between observed frequencies and population data, on the one hand, and (counterfactual) predictions concerning individual cases, on the other. Thus, our proposal is that IS generics that express empirical generalizations do so by pointing to a principled *causal* relation.

On this proposal it is also straightforward to account for the more liberal felicity, or acceptability, conditions of BP generics compared to IS generics: although BP generics express empirical generalizations that can already be determined by observing many instances and for which the frequency measure $\Delta^* P$ mostly seems appropriate,[19] this does not exclude the idea that the relative frequencies that make the use of BP generics acceptable should be explained by a causal analysis.

In fact, we will point out that making use of Pearl's causal models, a causal reading of generics correlates under many natural circumstances perfectly with the frequency measure $\Delta^* P$ that we have argued is appropriate for BP generics before.

4 CAUSAL READINGS OF GENERICS

4.1 *Causal models and the probability of causation*

We will use causal models (Pearl (2000)) (see also Spirtes *et al.* (2000)) to represent causal and counterfactual relationships. Causal models are closely related to Bayesian networks (Pearl, 1988). Such networks are able to represent probability functions in an efficient way, by making use of independencies. A Bayesian network consists of nodes that represent variables, and directed edges that are interpreted as representing direct dependencies between the variables. Special about causal models is that the edges are taken to represent 'causal mechanisms', i.e., stable and autonomous causal relationships, represented by equations. A causal model, $\mathcal{M}$, is a triple, $\langle U, V, F \rangle$, where U and V are disjoint sets of exogenous and endogenous variables, respectively, and F is a set of mappings. While the values of the *exogenous* variables U are determined by factors outside the model,

18 According to Popper, propensities are defined in terms of laws. With Skyrms (1980), we think it is more natural to go the other way around: propensities are probabilities/relative frequencies that are stable under conditioning with relevant background factors, and laws should be defined in terms of propensities. Still, we cannot quite follow Skyrms in all respects. Skyrms (1980) seeks to *define* causation in terms of (stable) probabilities. With Pearl (2000) we believe that any definition of causality in terms of probability will, in the end, be circular. With Pearl, we take causality to be primitive, and define causes of relative frequences in terms of it.

19 Not always though. First, 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. The sentence only seems appropriate if being a child of Mr. X somehow *explains why* one is a girl. Second, some people object that for 'taxonomic generics' like 'Men are mortal', $P(y|\neg x) = 0$, meaning that $\Delta^* P$ cannot be used for determining the appropriateness of such generics. Although we believe that there are ways around this problem, it is certainly the case that proposing a causal analysis for such generics is the most natural way to proceed.

the values of the *endogenous* variables V are determined by the values of variables in the model, i.e., by $U \cup V$. What U represents depends on the application of the causal model. An assignment to U, i.e, $U = u$, represents how things are that are not determined by the causal mechanisms in the causal model. Intuitively, U can normally best be thought of as a vector of variables, and $U = u$, i.e., $\vec{U} = \vec{u}$, as a vector of assignments to these variables. One can think of $\vec{U} = \vec{u}$ as the non-caused properties of a world/situation, or the non-caused properties of an individual, depending on whether the causal model represents causal relations between propositions, or between properties. An assignment of U uniquely determines the values of all the endogenous variables. Thus, if X is an endogenous variable, $X(u)$ gives the value of X if $U = u$. The value of $X(u)$ can, for instance, be the actual temperature, or the truth-value of p, if X stands for the issue whether p. But it can also stand for the length of an object, given the non-caused properties of such an object as set by $U = u$. The set of mappings F, finally, represent the mechanisms, or causal dependencies. Each $f_i \in F$ is a mapping which gives the value of V_i given the values of all other variables in $U \cup V$. More particularly, each function f_i can be written as an equation

(18) $v_i \quad = \quad f_i(pa_i, u_i),$

where pa_i denotes the values of the endogenous variables that are the *parents* of V_i, and where the u_i are the set of exogenous variables on which f_i depends.

What the parents of V_i are can be directly seen in the directed acyclic graph (DAG) associated with the causal model $\mathcal{M}$. Such a directed graph represents the casual dependencies between the members of V graphically. In such a graph, the members of V are represented by nodes, and if in $\mathcal{M}$ there exists a mapping f from pa_i (the members of V that are the parents of i) to V_i, this will correspond to directed edges from the nodes corresponding to the members of pa_i to V_i.

To illustrate, consider the following Firing Squad example. In this example we have the following (binary) random variables: A, B, C, D and U that are represented as nodes. A and B are riflemen (that shoot, a, b or not, $\neg a, \neg b$), C is the squad's captain (that orders c, or not $\neg c$) who waits for the court's order U (yes, or no). D, finally is the variable standing for death (or not) of the condemned prisoner. We assume that all variables are endogenous, except for U, which is exogenous. The laws can be stated as follows: $\{C := U, A := C, B := C, D := A \vee B\}$. Notice that we assume for simplicity that there are no separate exogenous variables related to A, B, C and D, representing the assumption that A and B are perfectly accurate and alert marksmen, and that they are, just as C, law-abiding. Moreover, we assume that D won't die from fear, a heart attack, or anything else but a shot. This situation, together with the causal relations, can be displayed by the following DAG.

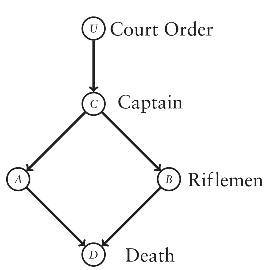

Notice that the value of the exogenous variable U determines the value of all other variables immediately: if $U = 1$, all endogenous variables will receive value 1 as well, and similarly when $U = 0$. But in many cases we are interested in what would be the case if something went different from how it actually went. Given that the prisoner is dead and that rifleman A shot him, we would like to know, for instance, whether the prisoner would still be dead had rifleman A not shot him. To answer such a counterfactual question, Pearl makes use of *interventions*.

For the representation of local actions, or hypothetical changes (i.e., interventions), Pearl (2000) makes use of *submodels*. A submodel $\mathcal{M}_x$ of $\mathcal{M}$ is the causal model $\langle U, V, F_x \rangle$, where x is a particular value of endogenous variable(s) X, and where F_x is just like F, except that all functions f_x that correspond to members of X are replaced by constant function $X = x$. Intuitively, $\mathcal{M}_x$ represents the minimal change from $\mathcal{M}$ required to make $X = x$ true for any $u \in U$. If X and Y are variables in V, the counterfactual 'The value that Y would have obtained, had X been x' is interpreted as denoting $Y_x(u)$. Intuitively, this will just be $f_y(x, u)$, in case X is the only parent of Y in the graphical model. In our Firing Squad example, the intervention of A can be represented by (i) breaking the causal link between the captain's order and the decision of A whether to shoot, and (ii) setting the value of A to 0.

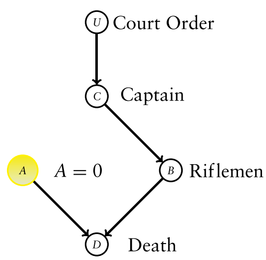

The counterfactual can now be handled as follows: First we reason backwards by making use of *abduction*: given that the prisoner is dead and that rifleman A actually shot him, we can conclude that the captain ordered a shooting, which means, in turn, that the court ordered this as well. Then, we make use of *intervention* (break the law $A := C$, and replace this with $A = 0$). Finally, we make use of *prediction*: what follows in the new imagined situation? Obviously, rifleman B is still ordered to shoot, so the prisoner would still be dead.

In this paper we will make use of a *probabilistic causal model*. This is a pair $\langle \mathcal{M}, P(u) \rangle$, where $\mathcal{M}$ is a causal model, and $P(u)$ is a probability function defined over the domain U.[20] Because each endogenous variable is a function of U, $P(u)$ completely defines the probability distribution over the endogenous variables as well. If Y is a variable in V, we will abbreviate $P(Y = y)$ from now on by $P(y)$. The latter is determined as follows:

$$(19) \quad P(y) \quad := \quad \sum_u P(u) \times \begin{cases} 1, & \text{if } Y(u) = y \\ 0, & \text{otherwise.} \end{cases}$$

20 In case there is more than one exogenous variable, $P(u)$ gives a probability to a vector of values to the variables in U.

Similarly, if X is also a variable in V, we will abbreviate $P(Y_x = y)$ by $P(y_x)$. The latter, which is also denoted by $P(y|do(x))$, is determined as follows:[21]

(20) $$P(y_x) \quad := \quad \sum_u P(u) \times \begin{cases} 1, & \text{if } Y_x(u) = y \\ 0, & \text{otherwise.} \end{cases}$$

Notice that if $P(y_x) > \frac{1}{2}$, this won't be due to the fact that in most situations where x were true, y would hold, but rather due to the fact that in most situations, y would hold after an intervention with x. Alternatively, if we think of x and y as properties, $Y_x(u) = y$ represents the 'fact' that after a (minimal) change of objects that satisfy properties $U = u$ such that it, or they, will become x (or will have property x), it will have property y. $P(y_x)$, then, measures, intuitively, the relative amount of objects that would be y after 'becoming' x due to a (hypothetical) intervention.

One of the most appealing features of calculating $P(y_x)$ as proposed by Pearl is that in this way we can also determine the probability $Y = y$ would have after an intervention that would make x true, if x is, in fact, not true. Thus, on Pearl's analysis one can easily determine $P(y_x \mid \neg x)$ as follows:[22]

(21) $$P(y_x \mid \neg x) \quad := \quad \frac{P(y_x, \neg x)}{P(\neg x)} \quad = \quad \sum_u P(u \mid \neg x) \times \begin{cases} 1, & \text{if } Y_x(u) = y \\ 0, & \text{otherwise.} \end{cases}$$

It is interesting to realize that to determine the counterfactual probability $P(y_x \mid \neg x)$ we will make use of (i) abduction, (ii) intervention, and (iii) prediction, just like in the above Firing Squad example. Abduction is used to determine from the assumed $\neg x$ what is the probability of u. Intervention is used to set the value of X to x, from which we calculate, by prediction, the probability of $Y = y$.

Pearl (2000, chapter 9) defines the '**probability of causal sufficiency**' of x to generate y, PS_x^y, as $P(y_x \mid \neg x, \neg y)$. We used bold font here, because it is in terms of this measure that we will propose to account for many of the above observations.

(22) $$PS_x^y \quad = \quad P(y_x \mid \neg x, \neg y) \quad = \quad \sum_u P(u \mid \neg x \wedge \neg y) \times \begin{cases} 1, & \text{if } Y_x(u) = y \\ 0, & \text{otherwise.} \end{cases}$$

Observe that in the Firing Squad example, $PS_a^d = P(d_a \mid \neg a, \neg d)$ will be 1. Indeed, a shooting of A suffices to cause the death of the prisoner.

4.2 *Relating the causal measures to the frequency-based ones*

We have suggested above that in contrast to BP generics, for their truth, or acceptability, IS generics require the existence of a 'principled connection', which we take to be a causal connection. But causal relations are famously hard to prove. Interestingly enough, however, under certain circumstances causal relations give rise to observable frequencies. This allows us to *test* whether a principled, or causal, connection exists, by looking at these frequencies.

Indeed, Pearl (2000, chapter 9) shows that the counterfactual probability measure $PS_x^y = P(y_x \mid \neg x, \neg y)$ comes down to a non-counterfactual probability measure under

21 For those who are not familiar with causal models, it might help to think of $P(y_x)$ as the probability of y after *imaging* P by x, as proposed by Lewis (1976), if X and Y correspond to variables associated with propositions.

22 For convenience, we abbreviate $P(x \wedge y)$ by $P(x, y)$.

some conditions. This is interesting for us, because it will allow us to relate the causal, and counterfactual, analysis of generics that we will discuss in the following subsection, with the (stable) frequency based analysis of BP generics discussed in section 2.

First, Pearl (2000, chapter 9, p. 294) shows that if a 'monotonicity' condition is assumed (the assumption that from a change from $\neg x$ to x, y cannot change into $\neg y$, or formally, $\forall u \in U : Y_x(u) \geq Y_{\neg x}(u)$), PS_x^y reduces to the following:

(23) $PS_x^y \quad = \quad P(y_x \mid \neg x, \neg y) \quad = \quad \frac{P(y_x)-P(y)}{P(\neg y,\neg x)}.$

It is easy to show that $\frac{P(y_x)-P(y)}{P(\neg y,\neg x)}$ is equal to $\frac{P(y_x)-P(y|\neg x)}{1-P(y|\neg x)}$, and thus that

(24) $PS_x^y \quad = \quad P(y_x \mid \neg x, \neg y) \quad = \quad \frac{P(y_x)-P(y|\neg x)}{1-P(y|\neg x)}.$

In section 2 we have argued that BP generics of the form 'xs are y' can be analyzed in terms of observed frequencies making use of the measure $\Delta^* P_x^y$. A crucial fact of this paper is that this latter measure is extremely closely related with PS_x^y. To get some more insight into how $PS_x^y = \frac{P(y_x)-P(y)}{P(\neg x,\neg y)}$ is related to $\Delta^* P_x^y$, one can show the following:[23]

(25) $PS_x^y \quad = \quad P(y_x \mid \neg x, \neg y) \quad = \quad \Delta^* P_x^y + \frac{P(y_x)-P(y|x)}{P(\neg y,\neg x)}.$

This means that PS_x^y is just $\Delta^* P_x^y$, corrected for by the second term to account for the case that $P(y_x) \neq P(y \mid x)$. Note that $P(y_x) \neq P(y \mid x)$, for instance, when x and y have a common cause.[24]

If one assumes in addition that the way Y would respond to x or $\neg x$ is independent of the actual value of x (what Pearl calls the 'exogeneity condition', which graphically means that there is no unblocked backdoor path from X to Y, which intuitively means that X and Y don't have a common cause), the causal influence of X on Y is reduced to conditional probability: $P(y_x) = P(y_x \mid x) = P(y \mid x)$. It is obvious that in that case PS_x^y comes down to the notion in terms of which we previously proposed to analyze generic sentences, i.e., $\Delta^* P_x^y$:

(26) $PS_x^y \quad = \quad P(y_x \mid \neg x, \neg y) \quad = \quad \frac{P(y|x)-P(y|\neg x)}{1-P(y|\neg x)} \quad = \quad \Delta^* P_x^y.$

23 Let us write $P(y)$ out as $P(y \mid x) \times P(x) + P(y \mid \neg x) \times (1 - P(x))$, or equivalently, $P(y) = P(y \mid x) \times (1 - P(\neg x)) + P(y \mid \neg x) \times P(\neg x)$. Now we can fill this latter formula into the above:

$$
\begin{aligned}
\text{(i)}\quad P(y_x \mid \neg x, \neg y) &= \frac{P(y_x)-P(y)}{P(\neg y,\neg x)} \\
&= \frac{P(y_x)-[P(y|x)\times(1-P(\neg x))+P(y|\neg x)\times P(\neg x)]}{P(\neg y,\neg x)} \\
&= \frac{P(y_x)-P(y|x)+P(y|x)\times P(\neg x)-P(y|\neg x)\times P(\neg x)]}{P(\neg y,\neg x)} \\
&= \frac{P(y|x)\times P(\neg x)-P(y|\neg x)\times P(\neg x)}{P(\neg y,\neg x)} + \frac{P(y_x)-P(y|x)}{P(\neg y,\neg x)} \\
&= \frac{P(y|x)\times P(\neg x)-P(y|\neg x)\times P(\neg x)}{P(\neg y|\neg x)\times P(\neg x)} + \frac{P(y_x)-P(y|x)}{P(\neg y,\neg x)} \\
&= \frac{P(y|x)-P(y|\neg x)}{1-P(y|\neg x)} + \frac{P(y_x)-P(y|x)}{P(\neg y,\neg x)} \\
&= \Delta^* P_x^y + \frac{P(y_x)-P(y|x)}{P(\neg y,\neg x)}.
\end{aligned}
$$

24 Notice that in case $P(y_x) \neq P(y \mid x)$, (i) $\Delta^* P_x^y$ might be high, although $P(y_x \mid \neg x, \neg y)$ is low, if $P(y_x) < P(y|x)$ e.g. if x and y have a common cause, and (ii) $\Delta^* P_x^y$ might be low, although $P(y_x \mid \neg x, \neg y)$ is high, if $P(y_x) > P(y|x)$. This has interesting consequences for a causal analysis of generics, but we won't go into that here.

It should be noted that Cheng (1997) derives $\Delta^* P_x^y$ as being the 'causal power' of x to induce y on similar, though still somewhat different, causal assumptions. Cheng (1997) assumes that the potential causes of y are independent of each other (the Noisy-OR assumption), and that $P(y_x)$ is independent of $P(x)$, meaning that $P(y_x \wedge x) = P(y_x) \times P(x)$.[25]

Finally, let us look at what happens if we assume not only that $P(y_x)$ is independent of $P(x)$, but also that y could only be caused by x. That is, $P(y|\neg x) = 0$. Now it immediately follows that not only $P(y_x)$, but also PS_x^y reduces to the conditional probability of y given x: $P(y|x)$. This is important, because according to Cohen (1999), it is under his 'absolute reading' enough for the generic to be true if $P(y|x)$ is high, irrespective of the value of $P(y|\neg x)$. So we see that Cohen's absolute reading of a generic comes out as a special case of our general analysis of (IS) generics, because under some conditions PS_x^y reduces to $P(y|x)$. This we take to be a pleasing feature of our analysis.

Let us illustrate our causal analysis with a non-generic example discussed in Pearl et al. (2016, p. 11-123), an example which highlights the use of the monotonicity assumption and the fact that our causal analysis is about *single case probabilities*. Consider the case of Mrs Smith, a woman with leukemia who had lumpectomy but not irradiation. Her tumor recurred after a year. She wonders whether she should regret her decision not to have gone through irradiation. Whether or not she should feel regret depends on the chance that *her* tumor would not have recurred, if she had gone through irradiation, which is clearly a case of *single case probability*. The chance that her tumor would not have recurred with irradiation, given that it did recur without irradiation is naturally measured by $PS_x^y = P(y_x|\neg x, \neg y)$, where '$x$' stands for 'going through irradiation', and '$\neg y$' for 'tumor recurred'. Of course, this measure is not, in general, estimable from observational or experimental data. As mentioned above, however, it is estimable in case we assume exogeneity and monotonicity. Exogeneity just means, as always, that $P(y_x) = P(y \mid x)$. Interestingly, for this case it is very clear what *monotonicity* means: irradiation cannot cause the recurrence of a tumor that was about to remit. With those two assumptions, whether Mrs Smith should regret her decision, which depends on whether the chance that *her* tumor would have recurred had she gone through irradiation, can be estimated by the value of the population data measured by $\Delta^* P_x^y$.

4.3 *IS and BP generics and causality*

In this section we will use causal models to give an analysis of IS generics. In our use of causal models we will think of x and y as properties of individuals, and take F to represent causal relations between properties.

Our (preliminary, see section 5) proposal now is obvious: IS generics of the form 'An x is y' are true, or felicitous,[26] iff PS_x^y is high. But this is only the case if any x by itself can, with high probability, be a *cause* of y. Or better, any individual with the property 'being an x' has a high chance of also having the property 'being a y' *in virtue of* having the former property. In this way, the intuition that there should be a 'principled connection' between x and y is accounted for. Moreover, on this proposal we implement the idea that IS generics are about

25 Recall that y_X denotes a proposition: the set of circumstances u for which it holds that $Y_X(u) = y$.

26 As mentioned before, we don't want to decide whether generics have truth conditions, or only acceptability conditions. In any case, we believe with Greenberg (2007) that the acceptability, or truth, of an IS generic is heavily dependent on the beliefs, stereotypes, norms, etc. of the speaker.

individual cases, because we don't first look at the frequencies, but rather at the causal structure that produces the frequencies. Many, if not most, BP generics are interpreted in the very same way. However, or so we would claim, a BP generic can already be true, or acceptable, if $\Delta^* P_x^y$ is high.[27,28]

Thus, for IS generics of the form 'An x is (a) y', it is a principled connection between x and y that counts. A principled connection that need not exist for the corresponding BP generic to be true, or acceptable. In this way, we can immediately explain the contrast between, for instance, the acceptable (9-a) versus the unacceptable (9-b), respectively (repeated here as (27-a) and (27-b)):[29]

(27) a. $\surd$ Musical comedies (plays) are popular.
b. $*$ A musical comedy (play) is (a) popular (play).

The fact that according to our analysis of IS generics, there should be a causal, or principled, connection between properties also explains why these generics are not good when they involve unnatural properties: causal connections between properties typically exist only when these properties are 'homogenous', and the unacceptable IS generics with unnatural properties discussed by Greenberg (2003, 2007) are not of that kind. Our causal account also explains why IS generics are not subject to scope ambiguities, which they would be on a quantificational analysis. Finally, our analysis can also explain why IS, but not BP generics are infelicitous, or unacceptable, with episodic predicates: causal relations are (taken to be) time independent, which is not the case for the IS generics (28-a) and (28-b):

(28) a. $*$ An earthquake is especially strong today.
b. $*$ An Italian restaurant is closed tonight.

Although we take it that for BP generics it can be just high $\Delta^* P_x^y$ that counts, we also think that in many cases this high value can be *explained* by our causal analysis: $\Delta^* P_x^y$ is high because PS_x^y is high.[30] Of course, for PS_x^y to come down to $\Delta^* P_x^y$ we have to make a few assumptions: either Pearl's (2000) exogeneity and monotonicity assumptions, or Cheng's (1997) noisy-OR (or noisy-AND) assumption. We take it that all of these are natural in our case.

27 We don't want to take a stand at this point on whether this means that BP generics are ambiguous between a causal and a non-causal reading, or whether their meaning is more underspecified.

28 We have argued in section 2 that for 'striking' BP generics like 'Sharks attack bathers', value, or utility, should play a role as well. Because the corresponding IS generics are typically taken to be awful, we think that in contrast to BP generics, for IS generics a low, but positive, value of $\Delta^* P_x^y$ cannot be compensated by a high measure of $V(y)$.

29 To be honest, how exactly to explain why (27-b) is bad, but (15-e) 'A refrigerator costs about 1000 euros' is ok eludes us so far.

30 Indeed, as mentioned by Haslanger (2013) many people are tempted to an essentialist reading of generics in general, and thus also of BP generics. For these people there will, on our analysis, be hardly any difference between the acceptability of an IS and a BP generic. Although we take this. on the one hand, to show that a causal analysis is natural, it also indicates, on the other hand, that it won't be easy to test whether IS and BP generics really give rise to different truth, or acceptability conditions. We feel, however, that people do make a distinction between the unacceptability of IS generics like (27-b), (28-a) and (28-b), on the one hand, versus the (potential) acceptability of their BP counterparts.

Exogeneity, or no confounding, formally means that $P(y_x) = P(y|x)$, which holds if high $P(y|x)$ is not due to a common cause. The monotonicity assumption says that x cannot prevent y to occur. Of course, we cannot assume this in general, but we can do so for any reasonable generic of the form 'xs are y' or 'An x is a y'. The noisy-OR assumption, finally, means that the potential causes for y, including x, are, or are taken to be, independent of each other. Notice that in case x is (taken to be) the only cause of y (perhaps due to essentialist' reasons), the latter assumption is met trivially. Recall that in that case PS_x^y comes down to $P(y|x)$.

Intuitively, what is going on in IS generics, or so we propose, is that for the interpretation of a (causal) generic of the form 'xs are y', or 'An x is y', we consider an individual or object that neither has property x nor y, and look, or imagine, how it would behave after we turned it, via an intervention, into an individual or object that has property x. (Or better, perhaps, that although it doesn't have property x, we perceive of it as having this property.) The claim is that whether the generic is true, or acceptable, depends on the causal structure. Perhaps more interestingly, the claim is that disagreement about the acceptability of the generic depends on disagreement about the underlying causal structure, rather than on disagreement of the statistics.[31]

We feel that making use of $PS_x^y = P(y_x \mid \neg x, \neg y)$ instead of, for instance, $P(y_x \mid \neg x)$ or $P(y_x)$, is important to explain away some seemingly obvious counterexamples to a probabilistic analysis of generics. Apparently, it is the case that about 80% of all chicken are female, due to the fact that, for economic reasons, most farmers gas male chicks immediately after birth. Still, a generic like 'A chicken is female' is completely unacceptable. But by calculating $PS_{chicken}^{female}$ one first looks only at objects that are neither chicken nor female. Then one estimates the chance with which the male object that is turned into, or imagined to be, a chicken is thereby turned into, or imagined to be, a female as well. But that seems absurd. A similar explanation might be given for why 'A person is more than 3 year old' is bad, although most people are over 3 years old. Again, to determine $PS_{people}^{>3\ years}$ you take some non-human animal that is not over 3 years old. But then, why should this animal be over 3 years old after an intervention?

In section 2 we have followed Cohen (1999) by restricting the 'domain' of the probability function to $\bigcup Alt(y)$ for the interpretation of a BP generic of the form 'xs are y'. We have shown that by doing so we can account for several examples that Leslie (2008) presented as fatal counterexamples to Cohen's probabilistic approach. To strengthen the relation between the 'statistical' analyses of section 2 with the present causal analysis, we should be able to let the alternative set $Alt(y)$ play a role now as well. Fortunately, it is easy to see how that works: instead of using $PS_x^y = P(y_x|\neg x, \neg y)$, we should use $P(y_x|\neg x, \neg y, \bigvee Alt(y))$. For a generic like 'A duck lays eggs' this means that PS_x^y doesn't measure the 'causal power' of ducks in general to lay eggs, but rather the causal power of *female* ducks to lay eggs, which is quite high, indeed.

31 Although this is perhaps best illustrated with BP generics like 'Muslims are terrorists', or 'Mexicans are rapists'.

5 'BACKTRACKING' AND 'ANALYTIC' IS GENERICS

Although we believe that the above analysis of IS generics is very natural, the way we have used causality until now is clearly too limited: there exist IS generics of the form 'An x is y' that are true and appropriate, even though x is not a physical cause of y. In this section we will discuss two such cases, and propose generalisations of our causal analysis to accommodate them as well.

Backtracking IS generics. Suppose we have a causal structure of the form $X \rightarrow Y$. It may well be that in such cases $\Delta^* P^x_y = \frac{P(x|y)-P(x|\neg y)}{1-P(x|\neg y)}$ has a high value, predicting that BP generics of the form 'Objects of type y are (generally) of type x' can be true, or acceptable, in such circumstances. On the causal analysis we have presented in the previous section, however, for the same sentence to be true or acceptable, it has to be the case that $P(x_y|\neg x, \neg y)$ is high. That, however, is impossible. In this type of causal structure intervention of Y doesn't influence the probability of x, and because now $\neg x$ is taken to hold, this means that $P(x_y|\neg x, \neg y) = 0$. Thus, an important empirical consequence of the analysis given above is that IS generics are predicted to be *asymmetric*.

Perhaps unfortunately, this seems to give rise to obviously false predictions. Suppose that all and only all hummingbirds have pink necks. Then it seems that both of the following are true or acceptable:

(29) a. A hummingbird has a pink neck. and
b. A bird with a pink neck is a hummingbird.

Also the following two generics both seem true or acceptable:[32]

(30) a. A tiger is striped. and
b. A cat-like animal that is striped is (normally) a tiger.

There is no denial of these facts, but it does seem to us that the latter sentences are true, or acceptable, for a different reason than the former ones. Intuitively, an analysis of generics (29-a) and (30-a) in terms of the analysis proposed in the previous section is, plausibly, correct, or so we think. The reason why the corresponding sentences (29-b) and (30-b) are also good is, we think, different. One difference is that (29-b) seems definitional, rather than empirical. But given examples like (30-b), that cannot be the general account of the difference.

We believe that we can provide a causal analysis not only of generics like (29-a) and (30-a), but also of generics like (29-b) and (30-b). But the causal analysis has to be *more general*. Although most generics of the form 'An x is y' are good because being an x causes feature y, others are good because being a y causes feature x. Or putting it in other words, while most IS generics have a *causal forward* reading, some generics appear to have an *evidential*, diagnostic, or *causal backward* reading. Importantly, having a causal backward reading is still in line with the general intuition that (descriptive) IS generics demand a 'principled connection', an intuition that we propose to cash out as the existence of a causal relation between x and y.

32 To be sure, we think that the generic (30-b) without the 'normally' is still acceptable. One reviewer disagrees, and mentions that 'it is well known that many sentences that are bad as IS generics are good when an overt quantificational adverb is introduced'. If this reviewer's view is correct, the generalised causal analysis proposed in this section as 'backtracking' IS generics can be ignored.

To give a more detailed account of the *causal backward* reading of generics, we can define the probability that, given y, y is due to x.[33] Denoting this measure by $P(x \rightsquigarrow y|\, y)$, it can be stated as follows:

(31) $P(x \rightsquigarrow y|\, y) \quad := \quad \frac{P(x) \times PS_x^y}{P(y)}$.[34]

We propose that in causal structures of the form $X \rightarrow Y$, IS generics of the form 'A y is an x' should be interpreted as true, or acceptable, just in case there is a relatively high probability that y is due to x, i.e., that $P(x \rightsquigarrow y|\, y)$ is high. Notice that in these causal structures $P(x \rightsquigarrow y|\, y)$ can be positive and high, even though $P(x_y|\neg x, \neg y) = 0 = PS_y^x$.

In the previous section we have shown that assuming exogeneity and monotonicity, $PS_x^y = \Delta^* P_x^y$. Remarkably enough, one can show that under these same conditions it also holds that $\Delta^* P_x^y = P(y \rightsquigarrow x|\, x)$, and thus that under exogeneity and monotonicity, $PS_x^y = \Delta^* P_x^y = P(y \rightsquigarrow x|\, x)$.[35] Thus, under both the causal forward *and* the causal backward 'reading', the truth, or acceptability, of an IS generic can be estimated by, or comes down to, the same frequency measure: $\Delta^* P_x^y$!

This shows that a generic of the form 'An x is (a) y' can be appropriate if $\Delta^* P_x^y$ is high under both causal structures $X \rightarrow Y$ and $Y \rightarrow X$. But in causal structure $X \rightarrow Y$, $\Delta^* P_x^y$ is high because $PS_x^y = P(y_x|\neg x, \neg y)$ is high, while in causal structure $Y \rightarrow X$, $\Delta^* P_x^y$ is high because $P(y \rightsquigarrow x|x)$ is high. In both cases, the truth, or acceptability, of the generic depends on there being a (strong) causal connection between x and y, measurable under natural circumstances by $\Delta^* P_x^y$, and given the causal structure it is always obvious how the causal connection should be cashed out: as high PS_x^y or as high $P(y \rightsquigarrow x|x)$.

33 As a reviewer remarks, Pearl (2000) rules out backtracking counterfactuals by fiat, though there is considerable empirical evidence that they are real. Interestingly enough, Dehghami et al. (2010) give strong arguments that backtracking counterfactuals are most acceptable when a lawlike relation is involved, even if it's a contingent one. In our analysis in the previous section we have followed Pearl using only 'forward' tracking counterfactuals. There are several ways to account for backtracking counterfactuals making use of causal models. In this paper we discuss a particular one that matches well with our analysis of BP generics.

34 This measure is taken from Cheng et al. (2007). Notice the similarity with Bayes' rule: $P(x|\, y) = \frac{P(x) \times P(y|x)}{P(y)}$.

35 To show that $\Delta^* P_x^y = P(y \rightsquigarrow x|\, x)$, note that by standard probabilistic reasoning it follows that $\Delta^* P_x^y = \frac{P(y|x) - P(y|\neg x)}{1 - P(y|\neg x)} = \frac{P(y|x) - P(y)}{P(\neg y \wedge \neg x)}$. By a similar reasoning it follows also that $\Delta^* P_y^x = \frac{P(x|y) - P(x)}{P(\neg y \wedge \neg x)}$. Crucially, one can also show that $P(y|x) - P(y) = \frac{P(y)}{P(x)} \times [P(x|y) - P(x)]$ (see below). From this and the previous results, it immediately follows that $\Delta^* P_x^y = \frac{P(y)}{P(x)} \times \Delta^* P_y^x$. Because $\Delta^* P_y^x = PS_y^x$ (under exogeneity and monotonicity) it follows that $\Delta^* P_x^y = \frac{P(y)}{P(x)} \times PS_y^x = \frac{P(y) \times PS_y^x}{P(x)} = P(y \rightsquigarrow x|x)$. Thus $PS_x^y = \Delta^* P_x^y = P(y \rightsquigarrow x|x)$. This leaves us to show that $P(y|x) - P(y) = \frac{P(y)}{P(x)} \times [P(x|y) - P(x)]$:

$$\begin{aligned} P(y|x) - P(y) &= \frac{P(x|y) \times P(y)}{P(x)} - \frac{P(x) \times P(y)}{P(x)} \\ &= \frac{1}{P(x)} \times [P(x|y) \times P(y)] - \frac{1}{P(x)} \times [P(x) \times P(y)], \\ &= \frac{1}{P(x)} \times [P(x|y) \times P(y) - P(x) \times P(y)] \\ &= \frac{1}{P(x)} \times P(y) \times [P(x|y) - P(x)]. \end{aligned}$$

We have seen in section 4 above that in special cases high PS_x^y comes down to high $P(y|\ x)$, i.e., the measure in terms of which Cohen (1999) accounts for the truth, or acceptability, of generics under their 'absolute' reading. Interestingly enough, in similar cases also $P(y \rightsquigarrow x|\ x)$ reduces to $P(y|x)$.[36]

Is the causal analysis so far general enough? Suppose y and x have a *common cause*. Would in those cases a generic of the form 'An x is y' be true, or acceptable? To give an example of a reviewer, suppose that someone (mistakenly) believes that smoking doesn't cause lung cancer, but that smoking and getting lung cancer are due to a common (genetic) cause (call it 'c'). Would somebody like that accept the generic 'A smoker gets lung cancer'? We don't know. However, we don't think this is impossible on our analysis, because there is now still a causal story to be told. Moreover, it seems clear how to measure the truth, or acceptability, of the generic 'An x is a y': by $P(c \rightsquigarrow x|x) \times PS_c^y$. Interestingly enough, one can calculate that if there are no other causal factors and if one makes again some natural assumptions, this measure comes down to $P(y|x)$.

Analytic IS generics. We have argued in section 3 against a 'definitional', or 'rules and regulation' account of IS generics: these accounts were not general enough to capture all possible cases. But a similar problem faces our account. Where the rules and regulation analysis could not account for IS generics that express empirical generalizations, we seem to be unable to account for IS generics that express, eh,.... rules and regulations.

(32) a. A disjunction is entailed by its disjuncts.
b. A maroon shirt is a red shirt.
c. A bachelor is an unmarried man.
d. A bishop moves diagonally (in chess).
e. A gentleman opens doors for ladies.
f. A person born in 1980 reaches the age of 40 in the year 2020.

Proposing a causal analysis for such examples seems far-fetched. And perhaps it is not needed. Notice that for all the above generics of the form 'An x is (a) y' it holds that $\Delta^* P_x^y = 1$, because $P(y|x) = 1$ and $P(y|\neg x) \neq 1$. Thus, we could possibly account for what is special about these above IS generics by claiming that it has to be the case that $\Delta^* P_x^y = 1$,[37] i.e., that $\Delta^* P_x^y$ must have its *maximal* value.

Perhaps. Still, we think that a causal-like analysis of the above examples is quite natural. Just as an ordinary generic like 'Rubber stretches' seems to be true *in virtue of* what it is to be rubber, (32-a), (32-b), (32-c), (32-d) (32-e) and (32-f) seem to be true or acceptable *in virtue of* what it means to be a disjunction, red, a bachelor, a bishop (in chess), a gentleman, or being born in 1980, respectively. A number of authors have accounted for 'in virtue of'- or 'because'-statements in terms of metaphysical grounding (e.g. Rosen (2010); Schnieder (2011); Correia & Schnieder (2012)).

36 To see this, note that $P(y \rightsquigarrow x|\ x) = \frac{P(y)\times PS_y^x}{P(x)}$. In the special cases where PS_y^x reduces to $P(x|\ y)$ it will thus hold that $\frac{P(y)\times PS_y^x}{P(x)}$ reduces to $\frac{P(y)\times P(x|y)}{P(x)}$. As a result, $P(y \rightsquigarrow x|\ x) = \frac{P(x\wedge y)}{P(x)} = P(y|x)$ in these cases.

37 Perhaps, rules and regulations are such that $P(y_x)$ is independent of $P(x)$ and $P(y)$. But we have seen that in that case PS_x^y and $P(y \rightsquigarrow x|x)$ reduce to $P(y|x)$.

Some have even suggested that grounding is something like 'metaphysical causation' (Schaffer (2016);Wilson (2017)) and that also grounding should be accounted for in terms of structural equation models, where intervention makes sense.

There exists a related way to 'liberate', or generalize, the notion of causation: turning it into 'influence'.[38] In section 4 we have modeled causation in terms of directed acyclic graphs (DAGs), where the direction of the arrows has a causal meaning. But causation can be thought of as only a special kind of *influence*, and a causal net can be thought of as just one type of *influence net*, which can be displayed by a graphical model as well. Williamson (2005) proposed to think of logic in terms of influence nets, and argues that — just like Bayesian nets for the representation and reasoning with probabilities — this has computational advantages. But once we do that, it seems that not only logical relations like (32-a), but also meaning relations as expressed by (32-b)–(32-f) can be analyzed in such an influence net. This idea is supported by work (e.g. De Raedt *et al.* (2016)) that shows the close relation between Bayesian networks and Logic Programming, two (still) popular frameworks for reasoning and knowledge representation. Bayesian networks — which one might think of as generalizations of causal networks — specify joint probability distributions over finite sets of random variables and consist of two components: (i) a qualitative or logical one that encodes the local influences among the random variables using a directed acyclic graph, and (ii) a quantitative one that encodes the probability densities over these local influences. A logical program consists of (Horn) clauses like $A \leftarrow B_1, \cdots B_n$, with a fact represented by $A \leftarrow \top$. A typical clause has a form like '$A \wedge B \leftarrow A, B$' or '$Mother - of(x, y) \leftarrow Parent - of(x, y) \wedge Female(x)$',[39] two rules that represent *conceptual* truths. The idea behind the connection between the two frameworks is that the logical component of Bayesian networks essentially corresponds to a logic program, where the random variables in the Bayesian network correspond to logical atoms. As it turns out, logic programming can be augmented with probabilities, which even strengthens the connection. Furthermore, the direct influence relation of Bayesian networks corresponds to the, so-called, 'immediate consequence operator', which is used to construct a (minimal) model of the program. Interestingly, the clauses of a logical program need not express conceptual relations: they typically can be given a *causal* interpretation as well. Intervention, too, can be made sense of for logical programs, for instance to account for counterfactuals (cf. Schulz (2014)).

To illustrate the idea, let us just look at (32-a) and (32-b). These correspond to $A \vee B \leftarrow A$ (and $A \vee B \leftarrow B$) and $R(x) \leftarrow M(x)$, respectively, which are stated, or derived, by the logical program that represents all conceptual truths. To determine whether, for instance, (32-b) is true, or acceptable, we have to check (on the forward 'reading') whether $PS^{R(x)}_{M(x)} = P(R(x)_{M(x)} \mid \neg R(x), \neg M(x))$ is high. Thus, we take an object (shirt) that is not maroon (and thus not red), and imagine, or make, it being maroon. Then we determine what the probability would be that it also would be red. But given the above logical program this is obvious: the probability will be 1. A similar reasoning shows that also (32-a) comes out true/acceptable, because obviously $P^{A \vee B}_{A} = P((A \vee B)_A \mid \neg A, \neg B) = 1$.

Whatever the exact details of metaphysical grounding and logic programming, we take this all as a strong indication that there exist natural general frameworks in which causal

38 In fact, Lewis (2000) does so as well, but in a quite different manner.

39 The variables are implicitly universally quantified.

and conceptual, or grounding, rules can be captured, and thought of in very similar, causal-like, ways. If so, examples like (32-b)–(32-f) ceases to be a problem for a causal(-like) analysis. But whether any of these suggestions stretches the notion of causation beyond recognition, we must leave to the reader.

6 CONCLUSION AND OUTLOOK

The goal of this paper was to account for the limited felicity, or acceptabilty, of IS generics compared to BP generics, on the one hand, while preserving the close similarity between the two types of generics, on the other. The second seems important, because claiming there to be a sharp distinction between the two types of generics seems unnatural. We proposed to capture the first goal by arguing that IS generics of the form 'An x is (a) y' are special, because they demand a causal(-like) relation between x and y. We have shown that by accounting for this hypothesis making use of causal models, the second goal can be satisfied as well.

We generalised our analysis to account for 'backtracking' generics, which we feel is very natural. And we had to stretch the initial idea behind the causal analysis somewhat to account for analytic' IS generics as well. Although we have argued in favor of the stretch, we can understand if this gives rise to some uncomfortable feeling. Moreover, it seems that our stretch of the notion of 'causation' is not yet wide enough: we have not yet discussed IS generics that express *social* or *artificial* kinds. As for social kinds, there seem to be IS generics involving social concepts that many people find acceptable, such as the following.[40]

(33) An American salutes the flag.

We think that (33) should receive a similar treatment as a generic like (32-e): it states what (according to the speaker) a 'real' American does, or should do. More problematic are IS generics involving artificial kinds expressing functional concepts:

(34) a. A religious ritual increases cohesion in society.
b. An OrangeCrusher 2000 crushes oranges.

Intuitively, the latter IS generic can be true even if no OrangeCrusher 2000 has ever been used, i.e. if no such machine has ever crushed oranges. The idea is, obviously, that this generic is nevertheless true, or acceptable, because it is *made for* that purpose. Although there is obviously a causal relation involved between what an OrgangeCrusher 2000 is, on the one hand, and what it does, on the other, a standard causal analysis won't work here. So there is still work to be done.

On the positive side, we think that our analysis of generics can be applied as well to some other phenomena. Causal models have long been used to account for counterfactuals, but an increasing number of authors have recently proposed to use them for the analysis of *indicative conditionals* as well: intervention is important for the analysis of many indicative conditionals. In contrast to most of these recent proposals, however, a causal treatment like

40 The example is provided to us by a reviewer. Obviously, there are many more IS generics involving social kinds that express racist beliefs that are acceptable for some members of our society. Disagreement on the truth, or the acceptability of such IS generics typically focusses on the underlying causal structure, or so we think.

ours would moreover account for the idea that there should be a connexion, i.e., a relevance relation, between antecedent and consequent (cf. van Rooij & Schulz (2019c)).[41] A second natural application of our framework involves *disposition ascriptions*. Such ascriptions have standardly been given an analysis in terms of counterfactual conditionals. Faced by some problems of such a treatment, Fara (2005) proposed an analysis in terms of generics. In van Rooij & Schulz (2019b) we argue that this proposal could be worked out more satisfactorily by basing it on our analysis of (IS) generics.

Funding

The work for this paper was supported by NWO grant 'From Learning to Meaning: A new approach to Generic Sentences and Implicit Biases' (dossiernummer 406.18.TW.007) from the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO) and by the grant 'The biased reality of online media - Using stereotypes to make media manipulation visible' (dossiernummer: 406.D1.19.059) of the same organisation.

References

Adams, E. W. (1998), '*A primer of probability logic*'. *CSLI Lecture Notes Number 68*. CSLI Publications. Stanford.

Asher, N. & M. Morreau (1995), 'What some generic sentences mean'. in Carlson and Pelletier, *The Generic Book*. Chicago: University of Chicago Press 300–39.

Burton-Roberts, N. (1976), 'On the generic indefinite article'. *Language* **52**: 427–48.

Burton-Roberts, N. (1977), 'Generic sentences and analyticity'. *Studies in Language*, **1**: 155–96.

Carlson, G. (1977), *Reference to Kinds in English*. Ph.D. dissertation, University of Massachusetts, Amherst.

Carlson, G. (1995), 'Truth-conditions of generic sentences: two contrasting views'. In Carlson G. & F. J. Pelletier (eds.),*The Generic Book*. University of Chicago Press. 224–37.

Carlson, G. (2009), 'Generics and concepts'. In Pelletier Francis J. (ed), *Kinds, Things and Stuff Mass Terms and Generics*. Oxford University Press. Oxford. 16–36.

Cartwright, N. (1989), *Nature's Capacities and their Measurement*. Oxford University Press. Oxford.

Cheng, P. W. (1997), 'From covariation to causation: a causal power theory'. *Psychological Review* **104**: 367–405.

Cheng, P. W. & K. Holyoak (1995), 'Adaptive systems as intuitive statisticians: causality, contingency, and prediction.' In J. Meyer & H. Roitblat (eds.), *Comparative Approaches to Cognition*. MIT Press. Cambridge, MA.

Cheng, P. W., L. R. Novick, M. Liljeholm & C. Ford (2007), 'Explaining Four Psychological Asymmetries in Causal Reasoning: Implications of Causal Assumptions for Coherence'. In J.K. Campbell et al. (eds.)*Topics in Contemporary Philosophy, vol. 4, Causation and Explanation*. MIT Press. Cambridge, MA. 1, 32.

Cohen, A. (1999), *Think Generic! The Meaning and Use of Generic Sentences*. CSLI Publications. Stanford.

Cohen, A. (2001), 'On the generic use of indefinite singulars'. *Journal of Semantics* **18**: 183–209.

41 In terms of probabilities, an intervention account would say that an indicative conditional is ok if $P(y_x)$ is high, while we would demand that PS^y_x is high. Importantly, if $P(y) = 1$, $P(y_x) = 1$ as well, while PS^y_x can still be low. As it turns out, our treatment of generics satisfies the standard axioms of connexion logic, a not well-known but still interesting kind of logic that wants to capture the intuition that there should be a connection between antecedent and consequent in order for a conditional to be true, or acceptable.

Cohen, A. (2004), 'Generics and mental representations'. *Linguistics and Philosophy* **27**: 529–56.

Correia, F. & B. Schnieder (eds.) (2012), *Metaphysical Grounding: Understanding the Structure of Reality*. Cambridge University Press. Cambridge.

Danks, D. (2014), *Unifying the Mind. Cognitive Representations as Graphical Models*. MIT Press. Cambridge, Massachusetts.

Dehghani, M., R. Iliev & S. Kaufmann (2012), 'Causal explanation and fact mutability in counterfactual reasoning'. *Mind and Language* **27**: 55–85.

De Raedt, L., K. Kersting, S. Natarajan & D. Poole (2016), *Statistical Relational Artificial Intelligence. Logic, Probability and Computation*. Morgan and Claypool Publishers. San Rafael.

Fara, M. (2005), 'Dispositions and habituals'. *Noûs* **39**: 43–82.

Greenberg, Y. (2003), *Manifestations of Genericity*. Routledge. London.

Greenberg, Y. (2007), 'Exceptions to generics: where vagueness, context dependence and modality interact'. *Journal of Semantics*, **24** 1–37.

Grimm, S. (2009), 'Number marking and individuation: a view from Dagaare'. M. Prinzhorn, V. Schmitt, and S. Zobel (eds.), *Proceedings of Sinn und Bedeutung 14. Vienna*, Austria. 168–84.

Harré, R. & E. Madden (1975), *Causal Powers: A Theory of Natural Necessity*. Basic Blackwell. Oxford.

Haslanger, S. (2010), 'Ideology, generics, and common ground'. In Charlotte Witt (ed.),*Feminist Metaphysics: Essays on the Ontology of Sex, Gender and the Self*. Springer. Dordrecht. 179–207.

Hoeltje, M. (2017), 'Generics and ways of being normal'. *Linguistics and Philosophy* **40**: 101–18.

Keil, F. (1989), *Concepts, Kinds, and Cognitive Development*. MIT Press. Cambridge, MA.

Keynes, J. M. (1952), *A Treatise on Probability*. Macmillan. London.

Krifka, M. (1987), *An Outline of Genericity*, [partly in collaboration with C. Gerstner]. University of Tübingen. SNS-Bericht.

Krifka, M. (2012), 'Definitional generics'. In Beyssade, C., Mari, A. & Del Prete, F. (eds.), *Genericity*. Oxford University Press. Oxford.

Krifka, M., Francis J. Pelletier, G. Carlson, A. ter Meulen, Chierchia G. & G. Link (1995), 'Genericity: an introduction'. In G. Carlson and F. Jeffry Pelletier (eds.),*The Generic Book*. University of Chicago Press. Chicago. 1–124.

Kripke, S. (1972/80), 'Naming and necessity'. In D. Davidson and G. Harman (eds.),*Semantics of Natural Language*. Synthese Library. Dordrecht. 253, 763–355, 769.

Lawler, J. M. (1973), 'Studies in English generic'. In *University of Michigan Papers in Linguistics 1: 1*. University of Michigan Press. Ann Arbor.

Leibniz, G. (1684), 'Meditations on knowledge, truth, and ideas'. In *republished in G.W. Leibniz: Philosophical Essays, edited and translated by R. Ariew & R. Garber*, vol. **1989**. Hackett Publishing Company. Indianapolis. 23–7.

Leslie, S. J. (2008), 'Generics: cognition and acquisition'. *The Philosophical Review* **117**: 1–47.

Leslie, S. J., S. Khemlani, S. Prasada & S. Glucksberg (2009), 'Conceptual and linguistic distinctions between singular and plural generics'. In *CogSci 2009 Proceedings*.

Lewis, D. (2000), 'Causation as influence'. *Journal of Philosophy* **97**: 182–97.

Nickel, B. (2009), 'Generics and the ways of normality'. *Linguistics and Philosophy* **31**: 629–48.

Oosterhof, A. (2008), *The Semantics of Generics in Dutch and Related Languages*. John Benjamins Publishing Company. Amsterdam.

Pearl, J. (2000), *Causality: Models, Reasoning and Inference*. Cambridge: Cambridge University Press.

Pearl, J., M. Glymour & N. Jewell (2016), *Causal Inference in Statistics. A Primer*. Wiley. Chichester.

Popper, K. (1959), 'The Propensity Interpretation of Probability'. In Korner & Price (eds.), *British Journal of the Philosophy of Science*. **10**: 25–42.

Prasada, S. & E. Dillingham (2006), 'Principled and statistical connections in common sense conception'. *Cognition* **99**: 73–112.

Putnam, H. (1975), 'The meaning of "meaning"'. In K. Gunderson (ed.), *Language, Mind and Knowledge*. University of Minnesota Press. Minneapolis, MN.

Rosen, G. (2010), 'Metaphysical dependence: grounding and reduction'. In B. Hale and A. Hoffmann (eds.), *Modality: Metaphysics, Logic, and Epistemology*. Oxford University Press. Oxford. 109–36.

Sankey, H. (1997), 'Induction and natural kinds'. *Principia: an International Journal of Epistemology* **1**: 239–54.

Shanks (1995), *The Psychology of Associative Learning*. Cambridge University Press. Cambridge.

Shep, M. C. (1958), 'Shall we count the living or the dead?' *New England Journal of Medicine* **259**: 1210–4.

Schaffer, J. (2016), 'Grounding in the image of causation'. *Philosophical Studies* **173**: 49–100.

Schnieder, B. (2011), 'A logic for "because"'. *The Review of Symbolic Logic* **4**: 445–65.

Schulz, K. (2014), 'Minimal models vs. logic programming: the case of counterfactual conditionals'. *Journal of Applied Non-Classical Logics* **24**: 153–68.

Skyrms, D. B. (1980), *Causal Necessity. A Pragmatic Investigation of the Necessity of Laws*. Yale University Press. New Haven and London.

Spirtes, P., C. Glymour & R. Scheines (2000), *Causation, Prediction, and Search*. MIT Press. Cambridge Mass.

Stalnaker, R. (1978), 'Assertion'. In P. Cole (ed.), *Syntax and Semantics, vol. 9, Pragmatics*. Academic Press, New York 315–32.

Tenenbaum, J. & T. Griffiths (2001), 'The rational basis of representativeness'. In J. Moore & K. Stenning (eds.), *Proceedings of the 23th Annual Conference of the Cognitive Science Society*. 1036–41.

Tenori, K., V. Crupi, N. Bonini & D. Osherson (2007), 'Comparison of confirmation measures'. *Cognition* **103**: 1007–19.

Tessler, M. & N. Goodman (2019), 'The language of generalization'. *Psychological Review* **126**: 395–436.

van Rooij, R. & K. Schulz (2019a), 'Generics and typicality: a bounded rationality approach'. *Linguistics and Philosophy*. doi: 10.1007/s10988-019-09265-8.

van Rooij, R. & K. Schulz (2019b), 'Natural kinds and dispositions: a causal analysis'. *Synthese*. doi: 10.1007/s11229-019-02184-y.

van Rooij, R. & K. Schulz (2019c), 'Conditionals, causality and conditional probability'. *Journal of Logic, Language and Information* **28**: 55–71. doi: 10.1007/s10849-018-9275-5.

Williamson, J. (2005), *Bayesian Nets and Causality*. Oxford University Press, Oxford. Philosophical and Computational Foundations.

Wilson, A. (2017), *'Metaphysical causation'*. Noûs. doi: https://doi.org/10.1111/nous.12190