ORIGINAL ARTICLE

Generics favor stability

Katherine Ritchie (University of California, Irvine) & Ny Vasil (California State University East Bay)

Correspondence

Department of Philosophy 85 Humanities Instructional Building Irvine, CA 92697

Email: ritchiek@uci.edu

Funding information

This material is based upon work supported by the U.S. National Science Foundation under Award Nos. 2315943 and 2315944. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Abstract: Humans seek to predict, explain, and control their environments. Generalizations—like those expressed by "children like candy" and "cigarettes cause cancer"—provide one resource to facilitate these tasks. We develop a proposal tying the acceptability judgments of generics to the psychological functions of generalizations. We argue that the notion of stability and a recognition of the varied scope of regularities provide resources for a unified account of an apparently diverse range of generics: causal and categorical, essentialist and structural. This unifying account coheres with the function of generalization in human psychology, and provides a new framework ripe for further investigation.

KEYWORDS

Generics; Generic generalizations; stability; contextual restriction; causation

1. INTRODUCTION

We are creatures who seek to predict, explain, and control our environments. Generalizations in thought and language provide one resource to facilitate these tasks. Generic generalizations (1-5), or *generics* for short, attribute properties to kinds or categories without relying on an overt quantifier like *every* or *most*.

1. Birds have wings.
2. Dogs bark.
3. Women are nurturing.
4. Black people face discrimination.
5. Immigrants hold poorly-paid jobs.

Generics have a number of properties that make them interesting to linguists, philosophers, and psychologists. They tolerate exceptions (not *all* dogs bark, but 2 is judged true). They can apparently be acceptable while holding of less than 1% of the instances of a kind (as in 6) all the way up to 100% of kind instances (as in 7).

6. Mosquitoes carry West Nile Virus.
7. Whales are mammals.

For these and other reasons, giving a semantic theory of generics is notoriously difficult. Our aim here isn't to offer a new semantic theory for generics. Rather, we develop a psychological account of an important feature driving people's acceptability judgments of generics rooted in what we take to be the psychological functions of generalizations–to facilitate prediction, explanation, and control. In doing so, we draw upon the notion of stability from the philosophy of causation (Lewis, 1986; Woodward, 2006, 2010, 2021). Roughly, a causal relationship is stable if it, as Woodward puts it, "would continue to hold under a 'large' range of changes in background circumstances or under background circumstances that are judged 'important' on the basis of subject matter specific considerations" (Woodward, 2010, p. 292). Stable causal relationships are not very sensitive to other variables having particular values (e.g., a particular temperature), which can be a good feature for scaffolding prediction, explanation, and intervention. Compared to a causal relationship

that is less stable with respect to the same set of background circumstances, the more stable relationship serves the psychological functions of prediction, explanation and intervention better, as relying on it equips one to make more accurate predictions, generate relevant explanations, and design more effective interventions across a wider range of circumstances, including novel environments.

We argue that acceptability judgments of generic claims are sensitive to the perceived stability of the relationships that they describe. Specifically, we propose that, all else equal, people view relationships perceived as relatively stable to be better candidates for generic generalizations, than those summarizing relationships seen as relatively unstable. Why? Because generics that express stable relationships will be better tools for accurately predicting and explaining across circumstances and designing effective interventions.

Our account is unifying in several ways. First, its focus on stability provides the basis for an account that unifies judgments about *causal generics* describing cause-effect relationships (8) and *categorical generics* describing relationships between kinds or categories and their features (1-5).

8. Smoking cigarettes causes cancer.

Work on causation has, obviously, focused on the former, while theories of generics have focused almost exclusively on the latter. A more unified account will provide insight into both.

Second, our account can accommodate generics with strikingly different explanatory commitments: essentialist (attributing a relationship to inherent kind nature, e.g. 1-3; Gelman, 2003; Leslie, 2017; Rhodes et al., 2012) and structural (attributing a relationship to stable external constraints acting on kind members in virtue of their positions in a broader structure, e.g. 4-5; Haslanger, 2016; Ritchie, 2019; Vasilyeva, Blanchard, & Lombrozo, 2018; Vasilyeva, Gopnik, & Lombrozo, 2018; Vasilyeva & Lombrozo, 2020; Ross, 2024).

Finally, in exploring these connections, our account reinforces the interdisciplinary bridge between linguistic, philosophical and psychological research, spanning the topics of causation, explanation, generalization, and language. The account thus opens possibilities for new interdisciplinary theoretical innovations and empirical research.

The paper is structured as follows. First, in Section 2, we present our account on which the acceptability judgements of causal and categorical generics are affected by perceived stability. We

call the view *Generics Favor Stability* or *GFS*. We provide support for GFS by drawing on empirical work in psychology and on theoretical work on generics in philosophy of language and linguistics (Section 3). Next, we consider ways regularities can hold broadly (i.e., globally) or only in particular societies, communities, workplaces, … (i.e., in restricted "bubbles") (Section 4). While more restricted regularities might appear unstable (holding in, but not out of a "bubble"), we suggest that meeting the aims of predicting, explaining, and intervening requires understanding and communicating about more than just global regularities; understanding and conveying regularities internal to "bubbles" is also important. So, we argue we should expect that generics can express contextually-restricted regularities that are stable within their more restricted scope. This accommodates structural generics (4-5) and calls for a more sophisticated picture of the psychological functions of generics. We flesh this out and summarize preliminary supporting evidence. Finally (Section 5) we wrap up and suggest ways our account opens new avenues for future research.

Before we proceed, three clarifications. First, in English, generics can be expressed using several forms. The most common and widely studied is the bare plural form, which includes a plural expression without an overt quantifier or article as in 1-7. Generics can also be expressed using indefinite singular expressions and definite singular expressions (9).

9. A whale / The whale is a mammal.

We restrict our attention to bare plural generics here, but we take it our account will also apply to other ways of expressing generics.[1] Second, we limit our focus to descriptive generics (like 1-8) and leave investigating normative generics (10) for future research.

10. Boys don't cry.

Finally, we frame our discussion in terms of the *acceptability* of generics, rather than their truth conditions. We use this term to refer to people's judgments about a claim being apt or appropriate

[1] In fact, it might be that judgments in favor of stability are even stronger for indefinite singular or definite singular forms, since these are sometimes taken to be definitional or to essential interpretation (Cohen, 2001), which would, as we explain below, require a high degree of stability.

to say in a context. This follows the precedent from philosophy of causation, as Lewis and Woodward claim that unstable causal claims are less acceptable or dispreferred, not that they are false. While we will make several claims about the semantics of generics–namely that they express relationships and, later, that these relationships can be contextually restricted–our aim is to offer a cognitive theory about how perceived stability affects generic acceptability judgments. Our account is compatible with many ways of understanding the semantics and pragmatics of generics.[2] Moreover, relying on a measure that is gradable rather than binary (a generic can be more or less acceptable), has practical benefits for carrying out empirical work.

2. GENERICS FAVOR STABILITY

On our account, relationships perceived as relatively stable are taken to be better candidates for generic generalizations than those perceived as relatively unstable. The account predicts that higher perceived stability of a relationship will lead to higher endorsement of a generic expressing that relationship (controlling for other factors, like average strength of a causal relationship or prevalence of an attributed property among kind members). We claim that this holds for both *causal* generics (*C causes E*) and *categorical* generics (*Ks F/ are Fs / have F*). In general, generics are more acceptable when they are perceived to express stable relationships. More stable generic generalizations support predicting, explaining, and controlling across a broader range of circumstances. To the extent people's perceptions of stability are well-calibrated, seeking stable generalizations will better underwrite these cognitive functions in novel circumstances. The breadth of generalization and exportability to novel cases are both recognized as achievements central to the role generalization plays in cognition (e.g., see Woodward (2010) on causation;

[2] Judgments about truth, acceptability, and assertability will often, we think, hang together. But, they can come apart. For instance, one might not accept the sentence "some of the dogs are brown" as a felicitous description of an array of four brown dogs. This could be due to a false quantity implicature, even if, upon reflection, one agrees that the sentence is technically true of the scene. While our account is compatible with views on which stability is part of a semantic theory of generics (see a related semantic view in Cohen (2004), discussed in Section 3 below), it is also compatible with non-semantic views that focus on pragmatic or cognitive features that are not taken to affect the truth of generics, but nevertheless affect acceptability. We bracket questions about the truth of generics and specific patterns of differentiation between acceptability and assertability here.

Lombrozo & Carey (2006) on explanation). And, we expect this to be reflected not just in cognition involving generalization, but in acceptability judgments of generics as well.

2.1 Causal Generics Favor Stability

Let's unpack the notion of stability. Stability (also called 'insensitivity' or 'invariance') was introduced by Lewis (1986) to account for differences in the acceptability of a range of causal claims that are, according to a counterfactual view of causation, all true. The notion has been further refined by Woodward (2006, 2010, 2021) as part of the interventionist view of causation. According to a simple version of a counterfactual dependence view of causation, C is a cause of E just in case if C had not occurred, E would not have occurred. Interventionist accounts of causation focus on interventions on C making a difference to E. While we won't take a stance on the correct view of causation, at least in some cases our causal judgments are well captured by these views. For instance, suppose that Billy threw a rock at a bottle and it shattered. If Billy hadn't thrown a rock at the bottle or there had been an intervention so the throw did not occur, the bottle wouldn't have shattered and we judge that the throwing of the rock caused the bottle to shatter. The causal claim concerning the effects of the rock throw is relatively stable as it would hold across many other circumstances (e.g., different historical events, different time of day, locations of many people and objects, external temperature …).

To illustrate the role of stability in a closely matched pair of examples, consider the following from Vasilyeva, Blanchard, and Lombrozo (2018):

> Case one: a medical journal reports an association between mutations in the BRCA1 gene and breast cancer.
>
> Case two: the same journal reports an equally strong association between mutations in the Gabrb1 gene and alcoholism.
>
> In both cases, the authors additionally report their results for participants of low versus high socioeconomic status (SES). The relationship between the BRCA1 gene and breast cancer holds with the same strength for both subgroups. In contrast, the relationship between the Gabrb1 gene and alcoholism holds strongly in the low-SES subgroup, but

nearly disappears in the high-SES subgroup.

Vasilyeva, Blanchard, and Lombrozo (2018) point out that even though the overall relationship between putative cause and effect is equally strong in these two hypothetical cases, the causal relationships differ in their stability across a background condition (SES levels): the former is relatively stable, the latter is relatively unstable.

This contrast aligns with Woodward's claim that a relationship is stable if it holds across a "'large' range of changes in background circumstances or under background circumstances that are judged 'important' on the basis of subject matter specific considerations" and, Woodward goes on to say, "to the extent that the relationship would be disrupted by changes in background circumstances, it is less stable" (Woodward, 2010, p. 292). We will return to considering what counts as a background condition and how to specify which of these are important in a moment. For now we spell out stability of a causal relationship as follows:

> **Stability (version 1):** the extent to which a causal relationship X→Y would hold in a variety of different background conditions.

The relationship between a Gabrb1 mutation and alcoholism is relatively unstable across SES conditions, and under GFS the corresponding causal generalization "mutations in the Gabrb1 gene cause alcoholism" will be judged as less acceptable than an equally strong but stable generalization "mutations in the BRCA1 gene cause breast cancer" (Vasilyeva, Blanchard, & Lombrozo, 2018). To quote Woodward once more, "we tend to regard causal claims that are highly sensitive [i.e., unstable] as defective, nonstandard, or at least importantly different from less sensitive [i.e., more stable] causal claims" (Woodward, 2006, p. 2).

The causal claim about the rock throw is about particular events, whereas the claims about gene mutations are about types of events. While both token (actual causation) claims and type (causal generic) claims can be assessed for stability of the underlying relationships, here we focus on the stability of generic generalizations, exploring the idea that causal generics tend to express relatively stable relationships. For instance, consider 8 (repeated below) and 11:

> 8. Smoking cigarettes causes cancer.

11. Rocks striking bottles cause them to shatter.

While there are exceptions to both, as not everyone who smokes cigarettes develops cancer and thrown rocks can strike bottles without breaking them, the relationships described are quite stable. They hold across locations–in South America, Africa, Europe–and across time periods. They hold across modal variations in nearby worlds involving many different changes. We take it that the perceived stability of the relationships reported contributes to the judgment that these causal generics are acceptable.

2.2 Causal *and* Categorical Generics Favor Stability

Perceived stability does not just affect the acceptability judgments of generic *causal* claims. Generics in general are more acceptable when they are perceived to be stable across changes in background conditions, including time, location, intrinsic and extrinsic features, and variation in category members across possible worlds or situations. This requires a reformulation of stability:

> **Stability (version 2):** the extent to which a relationship would hold across a variety of different background conditions.[3]

Stability is a modal notion – it is about what *would* hold in various conditions. Assessments of stability can be informed by a broad range of data: direct observations gathered across different background conditions within the actual world; testimony about such observations or their summaries; mental simulations of non-actual possibilities or counterfactuals, informed by prior beliefs about the relata and mechanisms connecting them, etc. For instance, one might consider whether a regularity holds of members of a species reliably across several habitats, and test this by sampling animals in these habitats. If a one finds that the property proves unstable across this way of partitioning actual members of the species, we predict this will decrease acceptability judgments of a generic reporting this relationship. Regardless of the source of the data, one is considering whether the relationship *would* hold across different background conditions.

[3] Here we use the broad category of relationships. We are focused on relationships between causes and their effects and relationships between categories and their properties (some of which are, we take it, causal as well). Other types of relationships could be assessed for stability as well.

The notion of background condition plays an important role in our account. Before offering our account, a note on terminology. We use "partition" to refer to all logically possible ways to chunk a domain along some variable or dimension. Partitioning a domain results in a set of cells (i.e., subsets) such that every member/instance is in exactly one of the cells. Partitions can involve conditions related to location, time, fur color, humidity… Domains can in principle be partitioned in ways that are not intuitively sensible or useful (e.g. we can partition along the time dimension into "dogs born on odd minutes" and "dogs born on even minutes"). Only a subset of the logically possible partitions will be eligible for assessing stability as only a subset are eligible background conditions, as we explain now.

Building on Woodward's notion of a background condition for causal relationships, we define background conditions as:

> A condition is a *background condition* relative to a causal relationship (C causes E) or a categorical relationship (Ks F/ Ks are Fs /Ks have F) being assessed just in case it is *not*:
>
> > (i) a condition "explicitly represented in the relationship" being assessed (e.g., being C, E, K or F) (Woodward, 2021, p. 241),
> > (ii) a condition causally between C and E (Woodward, 2021, p. 241),
> > (iii) a 'lower level' variable realizing C, E, K, or F, (Woodward, 2021, p. 241, fn17), or
> > (iv) another way of "unpacking" C, E, K, or F (e.g., into parts that jointly contribute to the relationship being assessed).

Condition (i) rules out C, E, K, and F as background conditions. Condition (ii) rules out mediators in causal chains between C and E. (iii) rules out cases like the temperature of a gas being realized by different combinations of kinetic energies of molecules (Woodward, 2021, p. 241, fn17). Finally, (iv) rules out other ways of "unpacking" a variable in the relationship being considered. For instance, Billy's throw of the rock could be broken down into parts like Billy winding up for the throw and Billy releasing the rock. That is, one way of unpacking a cause is into parts of the cause event which together constitute that event, and play complementary roles in generating the outcome. A kind might also be broken down into several component parts that constitute the kind, and play complementary roles in accounting for the attributed kind property. For example, the kind

'ducks' might be constituted by female and male ducks who play complementary roles when it comes to the kind's reproductive properties, such as "laying eggs". We develop this idea further in Section 3. For now it will suffice to say that in virtue of being ways of "unpacking" a variable that as a whole account for the attributed outcome/property, such partitions of causes and kinds do not count as background conditions across which a relationship is assessed for stability.

Eligible background conditions often include all sorts of features about the environment including location, time period, humidity, temperature, exposure to other elements, and so on. They can also include ways in which a relationship might hold across ways of partitioning a kind along features of its instances. For instance, in the case of categorical generics, partitions of kind members across sex, gender, race, age, or other features are eligible background conditions, unless ruled out as by one of the conditions (i)-(iv).

To illustrate, consider assessing 2, repeated below, for stability.

2. Dogs bark.

In fact, this regularity is stable across many times and places, many different breeds, actual and possible dogs, etc. And, we take it, people assume this. But, imagine that the relationship between dogs and barking was perceived to be unstable across a salient background condition; e.g., imagine someone came to find out it was extremely common in the Southern hemisphere, but rare in the Northern hemisphere (while keeping the average prevalence of barking in dogs fixed; i.e., in both stable and unstable scenarios, the overall proportion of barking dogs is constant; what varies is whether the barking dogs are evenly distributed across hemispheres (stable relationship) or concentrated in one hemisphere (unstable relationship)). On our account, controlling for other factors, highlighting instability across this background condition should decrease acceptability judgements of the corresponding generic (2).

Within the class of background conditions, we suggest that not all are taken to be equally relevant to assessing stability. While we cannot provide a full account of how background conditions are raised to salience or how people judge the (ir)relevance of a background condition, we offer three suggestions for future theoretical and empirical work. First, people seem to care more about stability across "normal, non-far-fetched variations in background conditions than [across] variations that are highly abnormal" (Woodward, 2021, p. 287). For example, in assessing

generalizations about cars, variations in road conditions (wet v. dry) will plausibly be taken to be relevant, whereas variations in abnormal events (e.g., being struck by lightning v. not) will plausibly be ruled to be irrelevant. Similarly, when assessing claims about birds, we suggest that background conditions involving accidents resulting in broken wings will often be deemed irrelevant due to their abnormality. In other words, some eligible backgrounds can be discounted if they are considered highly abnormal and, even if raised to salience, will not (significantly?) affect the acceptability judgments of a generic.

Second, we suggest that various psychological indices of how arbitrary or gerrymandered a partition is are likely pertinent to its perceived relevance when assessing a generic. Finally, context and features of the structure of a conversation might raise some dimensions and backgrounds to salience, making them particularly relevant. For instance, what shared information are we building on? What is the shared aim of inquiry or QUD? (Stalnaker, 1978, 1998; Roberts, 1996, 2012). Context can even make regularly discounted backgrounds salient and relevant for acceptability judgments (e.g., if we are talking about speeds nearing the speed of light, or an environment on the International Space Station).[4] We leave further exploration of these topics for future work.

GFS draws together two heretofore largely disparate areas of research–on the nature of causation/causal judgements and on generic representations. These are unified by what we take to be a cognitive preference, which we'll call the *stability preference*. This guides judgements about the acceptability of causal generics ('smoking causes cancer'), generics that might be represented as related to essence ('dogs bark'), as well as generics connected to social structures ('women make less than men', 'lawyers are obligated to defend their clients'; which we discuss in Section 4).

To clarify, our claim is not that all generics rely on causal reasoning or representations (although many might). Notice that we explicitly dropped any reference to causation in our version of stability (version 2). Rather, it focuses on relationships. These can, but need not, be causal relationships. For example, "Senators are thirty or older" is not true in virtue of senatorship having a causal impact on people's age. Similarly, in a discussion of chess, "Bishops move diagonally" is not a causal claim about how bishops move, but one about the rules specifying their moves. We

[4] We thank Isaac Wilhelm for pressing us on this issue and suggesting examples like these.

are also not claiming that generics like 'Ks are F' are true because Ks have a causal power to produce F or otherwise cause F (Rooij & Schulz, 2019, 2021). Rather, we developed a revised notion of stability from the causation literature to inform and develop our general claim that there is a connection between which generics we are apt to accept and our assumptions about their stability. This goes beyond causal circumstances or causal representations to also include normative and deontic properties like those related to obligations and permissions that might bear constitutive or definitional relationships to categories.[5]

Our account is meant to be inclusive (e.g., to apply to causal and categorical generics), but not exhaustive. That is, we do not claim that generic endorsement is only sensitive to perceived stability. Rather, GFS highlights stability as an important factor in generic endorsement, with an underappreciated scope of application. The account allows for other factors to affect and, possibly, override the effect of (in)stability (e.g. overall property prevalence within a class / causal strength, property distribution across mentioned or alternative kinds (Tessler & Goodman, 2019), and many other factors).

3. SUPPORT FOR GFS

Our account has empirical support and coheres with ideas central to theories of the semantics of generic generalizations. To those familiar with work on generics, there might seem to be something quite unsurprising (or, worse, old news) about the account we have sketched so far. In this section, we support its import by showing how it draws together a broad range of research in several disciplines with different theoretical commitments. Our account highlights a common preference that is supported by evidence from multiple sources and disciplines. We also show how GFS can handle a range of generics that are troubling for other accounts. For both of these reasons–the consilience of interdisciplinary evidence and better accommodating prima facie problem cases–we argue that GFS is an important component of a theory of generics and generalization.

Empirical work shows that people penalize both causal and explanatory generics when

[5] For example, Noyes et al. (2021) found that people judged deontic properties to hold of someone at the moment they became a member of an occupational category, whereas properties like skills were judged to hold after a longer period playing the role. They argue that the latter are more naturally interpreted as fitting with a causal representation whereas the former are more plausibly represented as intrinsically or constitutively connected to social kinds themself.

presented with evidence of instability across background conditions. For example, if a person learns that the effect of a bone-fortifying folate supplement varies greatly across people with different genotypes (i.e. is unstable across genotype A vs. genotype B), they are less likely to endorse a causal generic like "folate supplementation increases bone density", compared to a scenario where the efficacy of the supplement remains stable across genotypes. This holds even when the average causal efficacy of the supplement is the same for stable and unstable causes, suggesting that it is instability per se that drives the penalty in generic endorsement (Vasilyeva, Blanchard, & Lombrozo, 2018). Our ongoing work shows that the same instability penalty holds for categorical generics as well, and for generics describing relationships in both natural and social domains (Rouzbehani et al., 2025; Vasil & Ritchie, in prep.).

Research on generics and essentialism finds that at least some generics are taken to describe stable relationships (Gelman, 2003; Rhodes et al., 2012). Consider an example involving tigers: Suppose, as research suggests, that children and adults represent tigers as having an underlying causally potent essence. If the tiger essence is taken to cause features like having stripes, and individual tigers are all taken to have this essence, then regardless of where a tiger is, when it exists, or other features it has, it will be expected to have stripes. Essentialism easily gives rise to the assumption of a stable relationship between kinds and their properties. Each instance of a kind shares its essence and, so, each is likely to have properties that are caused by it, across a wide range of circumstances. Our argument that generics are more acceptable when they describe stable relationships fits neatly with the idea that generics are a good fit for describing essential properties.

In empirical studies, stability has been used as a marker of essentialism, exploiting the expectation that kind members with underlying essences have stable features even when environmental features are changed. In the switched-at-birth task, a creature of one kind is described as being adopted and raised by members of another species. For instance, a baby skunk might be adopted by dogs. Participants are then asked whether the baby animal will come to have properties of its biological kind (e.g., spraying an unpleasant smell) or its adoptive kind (e.g., barking). Children as young as 4 tend to think that for animal kinds, an individual will develop properties in line with its biological kind rather than those had by members of its adoptive kind (Gelman & Wellman, 1991; Gelman, 2003). This shows that some environmental features–in particular specific nurture contexts–are seen as less central to answering questions about induction than internal biological features. While the task involves a particular member of a kind being

switched at birth, rather than considering larger swaths of a kind in a different time, place, or other alternative circumstances, the switched-at-birth task taps into some of the same aspects as the notion of stability.

GFS also coheres with ideas from theories of generics. Many hold that generics convey non-accidental generalizations that are modally robust (e.g., Dahl, 1975; Carlson, 1995; Krifka et al., 1995; Pelletier & Asher, 1997; Greenberg, 2003, 2007; Nickel, 2008, 2016). While not developed as a semantic account, GFS taps into these ideas by positing that acceptability judgements are sensitive to whether an expressed relationship only holds in actual or very particular situations (with a precise temperature, history, specific individuals…), or across a broad range of counterfactual circumstances.

Stability also fits with observations about the interpretation of bare plural sentences based on classes of predicates. Carlson (1977) argues that the difference between what he calls *individual-level* and *stage-level predicates* patterns with whether a generic interpretation is available. When an individual-level predicate (e.g., 'have wings', 'human', 'intelligent') holds of something/someone, it holds for a long period, perhaps its entire existence. In contrast, stage-level predicates (e.g., 'be tired', 'be in New York') can hold of something for just a short stage of its existence. Notice that individual-level predicates are stable across a thing's existence (i.e., they seem to hold of the individual as a whole, hence the name) while stage-level predicates are not. Experimental work shows that properties like being tired are not deemed generalizable by children while properties like having four legs are (Gelman, 1988; Keil et al., 1998). GFS fits with the idea that generics and generalizations are preferred for stable rather than unstable properties.

Finally, stability is closely related to a feature of Cohen's (1999, 2004) probabilistic semantics of generics–the *homogeneity constraint*. The constraint requires that the truth of categorical generics–like those of the form 'Ks are F'–presupposes that F holds of roughly the same percentage of Ks across every segment of every psychologically salient partition of K. Partitions are ways of "chunking" the domain. For instance a kind domain might be partitioned into segments related to hair color, sex, subspecies, or genre/style. Suppose, for instance, a young/old partition is psychologically salient when considering 12.

12. Pandas eat bamboo.

Homogeneity requires that 12 is felicitous or true only if roughly the same percentage of young and old pandas eat bamboo. The homogeneity constraint gets onto the same pattern as stability, and we view Cohen's inclusion of this element in his semantic account of generics as an underappreciated insight.[6] Nevertheless, GFS brings additional advantages.

First, in connecting generic and causal judgments, the GFS account suggests a more general cognitive preference for conveying stable regularities, than what a semantics for generics straightforwardly suggests. On our view, the stability preference is explained via the cognitive function of judgements related to predicting, explaining, and manipulating. This preference might be built into the semantics of generics, as Cohen (1999, 2004) does. But the stability preference could also be taken to be part of a pragmatic account of generic judgements or an explanation of why we take some generics to be unacceptable, regardless of whether they really are false. On views like these, stability is not part of the semantics of generics, but is part of a complete cognitive explanation of human judgments. Our account is neutral between these options.

Second, stability includes a broader range of partitions and background conditions than the homogeneity constraint. Stability can be computed across such background conditions as variations across time, location, and environmental features like temperature (resulting from partitioning by time, location, temperature or other environmental characteristic, etc.). Cohen relies on features of specific conceptual representations–like ANIMAL or FURNITURE–to argue for what makes a partition eligible or ineligible. In contrast, we offer a more general account of background conditions that is not specific to any particular conceptual representations. Cohen argues that for many kinds (e.g., kinds picked out by basic level concepts like DOG or TIGER), partitions based on location or environment are ineligible (Cohen, 2004, pp. 545-546). However, variation in some features of dogs (say, their barking) across a partition into Northern and Southern hemispheres seems to be the very sort of instability that would lead to decreased endorsement of

[6] Cohen (2004) argues that the homogeneity constraint itself can be explained, in part, via features of cognition, namely via insights about how we cognize rules. His idea is that a more homogenous domain allows for easier learning of rules; generics are used to express default rules, so they should be easier to learn / deemed more acceptable when a domain is homogenous. We agree with Cohen that a cognitive explanation is an important element in explaining the acceptability of the use of generics. In this way, we see our account as intimately related to his. However, we focus on what we take to be the cognitive function of generics–as serving in prediction, explanation, and manipulations–and we find stability to be a good fit for exploring how these functions inform endorsement of generic generalizations.

a generic. Our account of background conditions can capture this.

Finally, our account of background conditions elegantly handles two potential counterexamples for Cohen's view. One central datum about generics is that they tolerate exceptions. Sentences like "Dogs bark" and "Smoking causes cancer" are acceptable (or true) even though not every dog barks and not everyone who smokes develops cancer. Yet, Leslie (2007) argues that this poses a problem for Cohen's homogeneity constraint (1999, 2004). When confronted with a sentence of the form 'Ks are F', one salient partition is what we call a *trivial partition.* A trivial partition segments Ks into those that are F and those that are not-F. Any trivial partition is one that involves extreme instability, holding of all Ks in one cell and no Ks in the other. So, the homogeneity constraint appears to falsely predict that generics cannot allow for exceptions. GFS avoids this problem. Condition (i) in the definition of a background condition states that any condition is ruled out as a background condition if it is explicitly represented in the relationship (i.e., this rules out C, E, K, or F). A trivial partition relies on precisely this sort of partition. It partitions based on one of the variables–e.g., F–when the relationship being assessed is that Ks are F. Trivial partitions reveal nothing about whether a relationship is stable across background conditions, as they are not eligible background conditions.

GFS also provides resources to explain the apparent instability (non-homogeneity) of so-called *minority characteristic generics* (Leslie, 2007, 2008).

13. Ducks lay eggs.
14. Sea turtles are long-lived.

While these generics are often accepted, they are not stable across sex-based (13) or developmental stage partitions (14). This poses a challenge to homogeneity, and prima facie, GFS as well. However, our account of background conditions again provides resources to explain this apparent instability.

Recall that condition (iv) rules a condition/way of partitioning ineligible to serve as a background condition if it involves "unpacking" C, E, K, or F into, e.g., parts when these jointly contribute to the relationship being assessed. If kinds are represented as having complex internal structures and these are represented as contributing to the relationship being assessed (e.g., F-

ing/being F), then the complementary components will not be perceived as eligible background conditions across which the stability of this relationship is computed.

Above we gave the example of a throwing event being "unpacked" into components like the wind up, flick of the wrist, and release. These are not background circumstances; they are parts of the cause itself. Together they contribute to the larger potentially causal relationship assessed. Stability is not assessed relative to variations in them.

Now, suppose that sexually reproducing biological kinds are represented as having male and female components that are complementary and integrated within reproduction. And suppose that biological kinds are represented as having both adult members and juvenile members that are integrated in the lifecycle of the kind. Our suggestion is that these are components of kind representations that are taken to contribute to properties related to reproduction and the life-cycle of the kind. We think this view is plausible (see also Leslie (2008)). If this is correct, a partition into male and female (or young and adult) instances is not a background condition across which stability of a property for which they jointly account is assessed. Such a purported partition would involve "unpacking" the kind itself (the K variable in the generic) into components of K, even though they jointly give rise to F. Preliminary empirical work (Beck et al., 2025) suggests that people's evaluations of minority characteristic generics are indeed sensitive to the cues of high vs. low interactivity among kind members, in ways consistent with our claim providing support that GFS can accommodate minority-characteristic generics.

In sum, GFS unifies research on causal and categorical generics, is supported by empirical work on causal judgements and on psychological essentialism, coheres with work on generics in philosophy of language and linguistics, and elegantly handles cases that have been problematic for related accounts. Next, we turn to considering the scope of regularities and the ways GFS can account for the acceptability of structural generics.

4. GLOBAL AND LOCAL REGULARITIES

As we just discussed, many theorists have argued that generics reinforce or elicit judgments that there is an underlying internal *essence* of a kind that gives rise to observable characteristics (e.g., Gelman, 2003, 2004; Leslie, 2017; Haslanger, 2011; Rhodes et al., 2012; Foster-Hanson et al.,

2016). Their research suggests that generics convey regularities that hold very broadly (i.e., globally). Others have argued that generics can be a conduit for *structural* thinking, conveying regularities that stem from stable external constraints acting on category members in particular societies or communities (Ritchie, 2019; Vasilyeva, Gopnik, & Lombrozo, 2018; Vasilyeva & Lombrozo, 2020). Successfully predicting, explaining, and intervening requires understanding and communicating about regularities that are restricted to socio-cultural or environmental "bubbles" as well as regularities that are global. For instance, recent research examined ways structural thinking might relate to proposing social interventions to remedy inequalities that hold in a particular society (Peretz-Lange et al., 2021; Amemiya et al., 2022).

We have argued that stability supports these functions, as more stable regularities will facilitate better prediction, explanation, and interventions across circumstances. Yet, structural generics appear to be unstable in important ways, as they do not hold across societies, locations, or time periods. And, people recognize this. For example, upon hearing a generic about a fictional immigrant group like "Borunians hold poorly-paid jobs" accompanied by a structural explanation attributing this association to local governmental regulations, people express an expectation that the regularity may be altered if a group member moves to another place (e.g. a Borunian who moves to another country may end up with a well-paying job; Vasilyeva & Lombrozo, 2020). Even young children recognize that generics can convey regularities that depend on and vary across contexts (Vasilyeva, Gopnik, & Lombrozo, 2018). When children receive a structural rather than internalist explanation of group properties, they increase their expectations that a child born to members of one group but brought up in an alternative context with members of another group will develop characteristics of the adoptive group (Vasil et al., in prep.).

So, GFS faces an apparent problem. While we've been arguing that a stability preference is a unifying feature guiding generic judgements, it seems structural generics are acceptable even when they are recognized to be unstable across societies, places, or time periods. To foreshadow our solution, we argue that regularities reported by generics come in different "sizes" or "scales". That is, some are global–holding across societies, time periods, …. Others hold only in local "bubbles", like a nation, community or workplace. In addition to summarizing broad generalizations, expressing local regularities would, in fact, be indispensable to serving the psychological functions of generics: by capturing such systematic (even if not global) patterns, generics would provide valuable guidance to situated agents navigating within "bubbles"

(Blanchard et al., 2018). Generics can, we argue, be implicitly contextually restricted when they report on a local regularity. In these cases, GFS only requires sensitivity to stability within a "bubble".

4.1 Contextually Restricted Generics

That generics can be contextually restricted is a familiar, albeit contentious, view (Nickel 2008, 2016; Sterken, 2015; Hesni, 2021). The most common variants of implicit contextual restriction in the literature involve an appeal to normality or to relevant members (Declerk, 1991; Pelletier & Asher, 1997; Nickel, 2008, 2016). For instance, one might argue that 15 is implicitly restricted to normal worlds or normal or relevant kind members, ruling out, e.g., pinnipeds that were in accidents resulting in flipper loss.

15. Pinnipeds have flippers.

These views involve restricting the domain to members of a kind (or worlds) that are normal or relevant or to relevant subkinds.[7] In contrast, we are focused on restricting to situational "bubbles". For example, we take 4 and 5 (repeated below) to express regularities that, unfortunately, hold within the "bubble" of American society today; these generics, we posit, are endorsed as contextually restricted claims.

4. Black people face discrimination.
5. Immigrants hold poorly-paid jobs.

Societies, nations, workplaces, schools, and time periods are some possible situations in which there can be robust, stable regularities that are restricted rather than global.

To further unpack the view, let's consider a generic uttered in the context of a discussion about social and political norms in the United States in the 1980s:

[7] See Sterken (2015) for arguments that generics are context-sensitive in their quantificational force and lexical restrictors.

16. Women have the right to choose.

Within this conversation, we take it 16 can be used to convey something like 16'.

16'. In the United States in the 1980s, women had the right to choose.

The restriction to the United States in the 1980s is implicit in 16. While in many places and time periods outside of this restriction women do not have the right to choose (e.g., much of America after SCOTUS overturned Roe v. Wade in 2022), within the contextual restriction in 16' the generic is stable. Women have the right to choose in 1981 and 1986, in New York, Minnesota, Florida, and so on. Further, if there were different women in America, they too would have the right to choose. Places outside of America or situations outside of the 1980s are irrelevant to assessing the stability of 16' and, given that we take that to be what is conveyed by 16 in the envisaged scenario, for the acceptability of 16 (for a contextual domain restriction account relying on similar examples, see Greenberg, 2007).

Notice that there is a closely related gloss of 16, given in 16'', that involves nominal modification (i.e., modification of the noun *women*) rather than sentential modification (i.e., modifying the entire sentence).

16''. Women in the United States in the 1980s had the right to choose.

While these are very similar, we take it generics are contextually restricted as in 16' as this can capture the widely held view that categorical generics are (at least very often) about kinds (Carlson, 1977; Nickel, 2016; Noyes & Keil, 2019; Ritchie & Knobe, 2020; Prasada & Dillingham, 2009; Prasada et al., 2013). An account which posits restricted noun phrases, like *women in the United States in the 1980s* in 16'', involves positing many generics which are not about kinds (or subkinds). Our account avoids this by restricting to situations. This results in only instances of a kind in the "bubble" being relevant to the acceptability of the generic without thereby requiring these to be generalizations about gerrymandered categories.

We close this section with several further instances of conversational contexts and

implicitly contextually restricted generic utterances (*a* sentences) and their explicitly contextually restricted variants (*b* sentences) (see Greenberg (2003, 2007) for further examples).

Conversational Context: Discussing a university where there is a norm for professors and lecturers to teach on different days of the week.

17. a. Professors teach on Tuesdays and Thursdays. Adjunct instructors teach on Mondays, Wednesdays, and Fridays.
    b. At this university, professors teach on Tuesdays and Thursdays. At this university, adjunct instructors teach on Mondays, Wednesdays, and Fridays.

Conversational Context: Discussing Saudi Arabia in 2017.

18. a. Women aren't allowed to drive.
    b. In Saudi Arabia at the time of the conversation, women aren't allowed to drive.

Conversational Context: Discussing Amish life in Lancaster County Pennsylvania[8]

19. a. Women's hair is never cut and is worn in a bun on the back of the head, concealed by a prayer covering.
    b. In Amish communities, women's hair is never cut and is worn in a bun on the back of the head, concealed by a prayer covering.

The generics in the *a* sentences are unstable when unrestricted. For instance, in many societies women cut their hair and are allowed to drive. In the provided conversational contexts, following our account, they convey restricted generics as in the *b* sentences. These are stable across times, places, and alternative members within the contextually restricted situations. Regularities can hold in local "bubbles" without holding globally and, we have argued, generics can be implicitly restricted to convey these.

4.2 Generics Don't Allow Implicit Contextual Restriction?

[8] Based on examples from https://www.discoverlancaster.com/amish/lifestyle/

While some have argued for restricting the domain of a generic to just normal or relevant kind members, or relevant subkinds, others have objected to this and the idea that generics allow for implicit contextual restrictions in the sense we just argued for has been widely disregarded (Dahl, 1975; Carlson, 1977; Krifka et al., 1995; Pelletier & Asher, 1997; Leslie, 2007; but see, Greenberg, 2003, 2007 and Sterken, 2015).

For instance, Krifka et al. (1995) argue that generics, unlike explicitly quantified claims like universal generalizations, can *never* be implicitly or covertly contextually restricted. They rely on a pair like the following (1995: 45).[9]

20. a. Every/Each professor wears a tie.
    b. At UCLA, every/each professor wears a tie.
21. a. Professors wear ties.
    b. At UCLA, professors wear ties.

They argue that while the explicitly quantified 20a can be used in a context to express 20b, there is no context in which 21a can be used to express 21b. Our cases 17-19 suggest that their judgment is incorrect. To draw out the situationally restricted interpretation we take to be available for 21a suppose that there is a dress code or informally adopted norm at UCLA that requires that professors dress formally. We're discussing UCLA and I utter 21a. It seems perfectly natural to hold that I have conveyed a restricted generalization like that expressed by 21b. Professors' attire at any other university is simply irrelevant to its acceptability. Moreover, our recent empirical findings suggest that adults, and to some extent children, do allow for generics to be implicitly contextually restricted (Ritchie et al., 2025; Vasil et al., under review).

We next turn to worries that go beyond acceptability judgments and that highlight the ways that stability is doing important work for structural generics as well. Positing that generics can be contextually restricted might make it seem like generics are used to express mere coincidental patterns rather than more robust regularities. In arguing against the possibility of contextual restriction, Krifka et al. state "A generic sentence states a lawlike regularity, and such regularity

[9] Krifka et al. (1995) use indefinite singular generics in their versions of 21. Since indefinite singular generics are more restrictive in several ways and given that we have focused on bare plural generics, we modify their examples here.

does not admit of contextual restriction" (Krifka et al., 1995, p. 45). If contextual restrictions allowed one to restrict to just the actual members of a kind or to the books on my shelf right now, it does indeed seem like generics would not be capturing lawlike regularities. For example, suppose that by matter of chance all of the actual past and present Supreme Court Justices have even social security numbers. Nevertheless an utterance of 22 is infelicitous (Dahl, 1975; Pelletier & Asher, 1997).

22. Supreme Court Justices have even social security numbers.

Or imagine that due to a series of tragic events every living koala has lost a paw. Even in this scenario, an utterance of 23 is not acceptable (Krifka et al., 1995; Pelletier & Asher, 1997).

23. Koalas have three paws.

A feature holding accidentally across all the actual members of a kind is not sufficient for the acceptability of a generic.

We agree that generics require non-accidental lawlike regularities; we think this is part of what stability is getting at.[10] But, this fails to show that generics are never contextually restricted. We take it that 22 and 23 are unacceptable because they are perceived to be unstable, not because people resist contextual restriction per se. If it is just a coincidence that every koala has lost a paw, this is not a regularity that is resilient to variations in instances of a kind and their contingent features. Many potential changes to the background conditions would result in koalas retaining all four paws. This lack of stability explains the unacceptability judgments about these generics. To draw this out a bit further, imagine a context in which one of these relationships is drastically more stable. For instance, imagine that there is a strictly enforced law requiring that koalas have a paw cut off at birth. We take it that in some conversations, 23 would then be a quite natural thing to say. This suggests that stability is driving judgments about the unacceptability of generics like 22

[10] At least for an important class of generics often called characterizing generics. Some generics might be acceptable (or true) just in virtue of a high proportion of the members of the category having the property or if a property is striking in some way (see, e.g., Leslie & Lerner, 2022)

and 23.[11]

On our account generics are preferable when stable. We have argued that implicit situational restrictions can make how things are at some other times or places irrelevant–it is just stability within the restricted "bubble" that matters–but stability holding to a relatively high degree *always* involves a regularity holding in alternative background circumstances of some sort or other. None of the restrictions we have posited rule out all alternatives so that only, say, the actual koalas alive now and the events that actually happened (all those tragic accidents!) are relevant when assessing the generic.

5. CONCLUSION AND FUTURE DIRECTIONS

GFS posits a unifying psychological proposal–generics are favored when they are perceived to convey stable regularities–tied to the functions of generalization. We argued that the stability preference underlies both causal and categorical generics, regardless of whether they convey generalizations taken to be essentialist or structural. Our account accommodates the assumption of context-independence commonly associated with essentialist generics, while also capturing the ways structural generics report societal regularities based on laws, institutional practices, cultural norms, and the like, that are systematic yet not necessarily global (i.e., hold reliably within, but not necessarily across, contextual "bubbles"). Recognizing that regularities come in different sizes–some global, others restricted to smaller "bubbles"–set the stage for understanding how the acceptability of both essentialist and structural generics is sensitive to stability.

As we've emphasized, a stability preference also coheres with what we take to be the functions of generic judgments and representations. We are creatures that seek to predict, explain, and exert control across a wide range of locations, times, and circumstances. If a generalization is stable, it will provide better predictive and explanatory guidance and better serve our efforts to

[11] This case is akin to Nickel's (2016) discussion of ways in which conversational topic can affect the felicity (or truth) of a generic. He argues that the differences between breed standards for Dobermans and the features Dobermans have without body shaping, can account for the acceptability of (i) in discussions of evolutionary biology and (ii) in the context of discussing dog breeding or features of dogs at a dog show.

(i) Dobermans have floppy ears.

(ii) Dobermans have pointy ears.

control and shape our environment. Thus, we have argued people prefer generics that they perceive to express stable relationships.

Our account suggests additional questions to be explored: To what extent does the stability preference hold across causal and categorical generics? Does it hold equally across domains (e.g., animal and social) or is it stronger in some? To what extent does perceived stability affect our judgments about different types of generics, including majority and striking property generics? (It seems plausible that generics attributing striking properties, like "sharks attack people" and "ticks carry Lyme disease" are, in virtue of highlighting a rare but highly dangerous attribute, less sensitive or even completely insensitive to its stability). Likewise, can perceived or intended stability shape endorsement of normative generics, like "boys don't cry"? (Computing the stability of normative claims poses interesting theoretical and practical challenges - what are the relevant backgrounds? - such that stability may ultimately not be computable for this class of generics, rendering it irrelevant). Does stability affect judgments of generic acceptability, assertability and truth in the same way, or do they come apart? Does instability along some dimensions (e.g., counterfactual vs. alternative geographic locations) matter more (see also, Cohen (2004))? To what extent do the stability preference and situational restrictions vary across development?

The account also leaves open the question about the extent to which both causal and categorical generics are implicitly restricted across various domains (e.g., social, natural). In arguing for implicit contextual restriction we focused on social cases related to laws, norms, causal regularities, and other practices within social "bubbles" like countries and universities. Implicit situational restrictions might be much more common in the social domain than in others. If this is borne out empirically, we suggest it should be explained not by differences in the nature of these generics, but by differences in how people understand and communicate about regularities across these domains.

Finally, there's the question whether our account might extend beyond generics. We focused on generics, and argued that generics contribute to the psychological functions of predicting and controlling agents' environments; the sensitivity of generic acceptability to the perceived stability of causal and categorical relationships supports more robust predictions and interventions. Other linguistic forms might also contribute to these functions. We speculate that to the extent other linguistic forms serve the same functions, they are likely to also be sensitive to stability – but establishing this is a direction for future work. Finding ways to not only document

the stability preference but to establish its origin in each case is not a trivial matter. For example, token causal and explanatory claims also appear to be sensitive to the stability of the relationships they express. This includes the counterfactual stability of a token causal relationship (whether the relationship would continue to hold under variations in circumstance; e.g., would the rock throw still break the window if the weather were warmer, etc.; Woodward, 2006), and the stability of type-level relationships across actual contexts (e.g., people's evaluations of a claim that consuming a certain food caused a particular symptom varied with the stability of that causal relationship at the type-level, when participants were uncertain in which context the token instance took place; Vasilyeva, Blanchard, & Lombrozo, 2018). This shared stability preference could signal that token-level causal-explanatory claims serve the same purposes as type-level causal-explanatory generics (contributing to prediction and control, albeit perhaps less directly), or that token claims are treated as derivative from the corresponding generics and inherit their features.[12] These are all important questions, calling for a systematic empirical investigation of which expressions in natural language are and are not sensitive to stability, the degree to which they are, and how these relate to the stability of generics.

Overall, the Generics Favor Stability account unifies important insights from philosophy of science with work in philosophy of language, linguistics, and psychology. It can account for a wide range of generics, while being supported by a plausible account of the function of generalization in human psychology. It yields new constraints on future theorizing about generics: any adequate account will need to have the explanatory resources to accommodate stability effects. And, it provides a framework ripe for further interdisciplinary empirical and theoretical investigation.

Acknowledgements

We presented versions of the ideas in this paper at the Society for Philosophy and Psychology, the European Society for Philosophy and Psychology, Social Ontology, Cognitive Science Society, Cognitive Development Society, Society for Research on Child Development, Society for Personality and Social Psychology, Western Psychological Association, Spartan Psychological

[12] We thank an anonymous referee for suggesting we consider these questions further and for suggesting possible connections between token causal claims and generic causal claims.

Association Research Conference (SPARC), the Social Identities and Cognition in the Desert workshop, the Generic Workshop at UC Irvine, Genericity, Stability and Structural Interactions Workshop at California State University East Bay, the Intuitive Theories of Social Structures and Social Change workshop at Princeton University, the Duke Causation Workshop, the UCLA Center for Behavior, Evolution & Culture Talk Series, the UCSD Developmental Brownbag Talk Series, the University of Idaho Interdisciplinary Data Science and the Department of Politics and Philosophy Seminar Series, Annual UC Davis Symposium on Language Research, and colloquia at the National University of Singapore (NUS), Northwestern University, UCLA, and University of Nevada Reno. We thank audiences at these talks and poster presentations for their feedback. In particular, we thank Sam Cumming, Samia Hesni, Josh Knobe, Sandeep Prasada, Lauren Ross, Henry Schiller, and Isaac Wilhelm for comments and discussions that significantly improved the paper. We also thank the editors and two anonymous referees for their helpful comments.

References

Amemiya, J., Mortenson, E., Heyman, G. D., & Walker, C. M. (2022). Thinking structurally: A cognitive framework for understanding how people attribute inequality to structural causes. *Perspectives on Psychological Science*, *18*(2), 259-274. https://doi.org/10.1177/17456916221093593

Beck, J., Li, Y., Dutta, R., Zhu, T.A., Pederneschi, A., Curiel, A., Graver-Dowd, C., Greyson, I.N., Ultreras-Castro, K., Ritchie, K., & Vasil, N. (July-August, 2025). Do ducks lay eggs? How interactivity shapes generic generalizations. Poster presented at the 47th Annual Meeting of the Cognitive Science Society. San Francisco, CA.

Blanchard, T., Vasilyeva, N., & Lombrozo, T. (2018). Stability, breadth and guidance. *Philosophical Studies*, *175*(9), 2263-2283.

Carlson, G. (1977). *Reference to kinds in English*. Doctoral dissertation. University of Massachusetts, Amherst.

Carlson, G. (1995). Truth-conditions of generic sentences: Two contrasting views. In G. Carlson & F. Pelletier (eds.). *The generic book* (pp. 224-237) Chicago: University of Chicago Press.

Cohen, A. (1999). Think Generic!: The meaning and use of generic sentences. Stanford: CSLI.

Cohen, A. (2001). On the generic use of indefinite singulars. *Journal of Semantics*, *18*(3), 183–209. https://doi.org/10.1093/jos/18.3.183

Cohen, A. (2004). Generics and mental representation. *Linguistics and Philosophy*, *27*(5), 529–556.

Dahl, O. (1975). On generics. In E.L. Keenan (Ed.) *Formal semantics of natural language* (pp. 99-111) Cambridge: Cambridge University Press.

Declerk, R. (1991). The origins of genericity. *Linguistics*, *29*(1), 79–102.

Foster-Hanson, E., Leslie, S. J., & Rhodes, M. (2016). How does generic language elicit essentialist beliefs? In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 1541–1546). Cognitive Science Society.

Gelman, S.A. (1988). The development of induction within natural kind and artifact categories, *Cognitive Psychology*, *20*(1), 65-95, https://doi.org/10.1016/0010-0285(88)90025-4.

Gelman, S.A. (2003). The essential child: Origins of essentialism in everyday thought. Oxford University Press.

Gelman, S.A. (2004). Psychological essentialism in children. *Trends in Cognitive Sciences*, *8*(9), 404-409.

Gelman, S. A., & Wellman, H. M. (1991). Insides and essence: Early understandings of the non-obvious. *Cognition, 38*(3), 213–244. https://doi.org/10.1016/0010-0277(91)90007-Q

Greenberg, Y. (2003). *Manifestations of genericity*. New York: Routledge.

Greenberg, Y. (2007). Exceptions to generics: Where vagueness, context dependence and modality interact. *Journal of Semantics*, *24*(2), 131-167.

Haslanger, S. (2010). Ideology, generics, and common ground. In C. Witt (Ed.), *Feminist metaphysics: Explorations on the ontology of sex, gender and the self* (pp. 179-207). Springer. https://doi.org/10.1007/978-90-481-3783-1_11

Haslanger, S. (2016). What is a (social) structural explanation? *Philosophical Studies, 173*(1), 113-130.

Hesni, S. (2021). Generics as instructions. *Synthese*, *199*, 12587-12602. https://doi.org/10.1007/s11229-021-03344-9

Keil, F.C., Smith, W.C., Simons, D.J., Levin, D.T. (1998) Two dogmas of conceptual empiricism: Implications for hybrid models of the structure of knowledge, *Cognition*, *65*, 103-135.

Krifka, M., Pelletier, F. J., Carlson, G. N., ter Meulen, A., Link, G., & Chierchia, G. (1995). Genericity: An introduction. In G. N. Carlson, & F. J. Pelletier (Eds.), *The generic book* (pp. 1-124). The University of Chicago Press.

Leslie, S.J. (2007). Generics and the structure of the mind. *Philosophical Perspectives*, *21*(1), 375–403.

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

Leslie, S.J. (2017). The original sin of cognition: Fear, prejudice, and generalization. *The Journal of Philosophy*, *114*(8), 1–29.

Leslie, S.J., & Lerner, A., (2022). Generic Generalizations. In E.N. Zalta & U. Nodelman (Eds.) *The Stanford Encyclopedia of Philosophy* (Fall 2022 Edition), URL = <https://plato.stanford.edu/archives/fall2022/entries/generics/>.

Lewis, D. (1975). Adverbs of quantification. In E. L. Keenan (Ed.), *Formal semantics of natural language* (pp. 3-15). Cambridge University Press.

Lewis, D. (1986). Causation and postscripts to causation. *Philosophical Papers*, *2*, 172-213.

Lombrozo, T., & Carey, S. (2006). Functional explanation and the function of explanation. *Cognition, 99*(2), 167-204.

Nickel, B. (2008). Generics and the ways of normality. *Linguistics and Philosophy*, *31*, 629-648.

Nickel, B. (2016). Between logic and the world: An integrated theory of generics, Oxford: Oxford University Press.

Noyes, A., Dunham, Y., Keil, F.C., & Ritchie, K. (2021). Evidence for multiple sources of inductive potential: occupations and their relations to social institutions. *Cognitive Psychology*, *130*, 101422. https://doi.org/10.1016/j.cogpsych.2021.101422

Noyes, A. & Keil, F. C. (2019). Generics designate kinds but not always essences. *Proceedings of the National Academy of Sciences*, *116*(41), 20354-20359. https://doi.org/10.1073/pnas.1900105116

Pelletier, F. & Asher, N. (1997). Generics and defaults. In J. van Benthem & A. ter Meulen (Eds.), *Handbook of logic and language* (pp. 1125-1177). MIT Press.

Peretz-Lange, R., Perry, J., & Muentener, P. (2021). Developmental shifts toward structural explanations and interventions for social status disparities. *Cognitive Development*, *58*, 101042. https://doi.org/10.1016/j.cogdev.2021.101042

Prasada, S., & Dillingham, E. M. (2009). Representation of principled connections: A window onto the formal aspect of common sense conception. *Cognitive Science*, *33*(3), 401–448. https://doi.org/10.1111/j.1551-6709.2009.01018.x

Prasada, S., Khemlani, S., Leslie, S.J. & Glucksberg, S. (2013). Conceptual distinctions amongst generics. *Cognition*, *126*(3), 405–422.

Rhodes, M., Leslie, S. J., & Tworek, C. M. (2012). Cultural transmission of social essentialism. *Proceedings of the National Academy of Sciences*, *109*(34), 13526–13531. https://doi.org/10.1073/pnas.1208951109

Ritchie, K. (2019). Should we use racial and gender generics? *Thought: A Journal of Philosophy*, *8*(1), 33-41. https://doi.org/10.1002/tht3.402

Ritchie, K. & Knobe, J. (2020). Kindhood and essentialism: Evidence from language. In M. Rhodes (Ed.), *Advances in Child Development and Behavior*, *59* (pp. 133-164). https://doi.org/10.1016/bs.acdb.2020.05.005

Ritchie, K., Curiel, A., Zhu, T.A., Pederneschi, A., & Vasil, N. (2025). The scope of generic generalizations: Developmental changes in judgments about contextually-restricted generics. Paper published in the *Proceedings of the 47th Annual Meeting of the Cognitive Science Society*. San Francisco, CA.

Roberts, C. (1996). Information structure in discourse: towards an integrated formal theory of pragmatics. In J.-H. Yoon & A. Kathol (Eds.), *Papers in semantics* (*working papers in linguistics* 49). The Ohio State University.

Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics, 5*(6), 1-69.

Ross, L. (2024). What is social structural explanation? A causal account. *Noûs, 58, 163–179.* https://doi.org/10.1111/nous.12446

Rouzbehani, B., Caoile, J., Campos, G., Khurelkhuu, S., Curiel, A., Ritchie, K., & Vasil, N. (April, 2025). The instability penalty in endorsement of causal and categorical generics. Poster presented at the *California State University East Bay Opportunities for Student Scholarships, Creative Activities and Research (OSCAR) Annual Research Expo*. Hayward, CA.

Stalnaker, R. (1978). Assertion. In P. Cole (Ed.) *Syntax and semantics 9: Pragmatics* (pp. 315-332). New York Academic Press.

Stalnaker, R. (1998). On the representation of context. *Journal of Logic, Language and Information, 7*(1), 3-19.

Sterken, R. K. (2015). Generics in context. *Philosophers' Imprint*, 15(21), 1-30. http://hdl.handle.net/2027/spo.3521354.0015.021

Tessler, M. H., & Goodman, N. D. (2019). The language of generalization. *Psychological Review*, *126*(3), 395–436. https://doi.org/10.1037/rev0000142

van Rooij, R., & Schulz, K. (2019). Generic sentences: Representativeness or causality?. *Proceedings of Sinn und Bedeutung*, *23*(2), 409-426.

van Rooij, R., & Schulz, K. (2021). A causal power semantics for generic sentences. *Topoi*, *40*, 131–146. https://doi.org/10.1007/s11245-019-09663-4

Vasil, N., Curiel, A., Zhu, T. A., Pederneschi, A., & Ritchie, K. Living in a bubble: How children and adults use generics to talk about social kinds in unique sociocultural contexts. (Manuscript under review)

Vasil, N., & Ritchie, K. (in prep.). Generic generalizations about social groups and natural kinds: the role of stability.

Vasilyeva, N., & Lombrozo, T. (2020). Structural thinking about social categories: Evidence from formal explanations, generics, and generalization. *Cognition*, *204*, 104383. doi.org/10.1016/j.cognition.2020.104383

Vasilyeva, N., Gopnik, A., & Lombrozo, T. (2018). The development of structural thinking about social categories. *Developmental Psychology,* *54*(9), 1735-1744. https://doi.org/10.1037/dev0000555

Vasilyeva, N., Blanchard, T., & Lombrozo, T. (2018). Stable causal relationships are better causal relationships. *Cognitive Science*, *42*(4), 1265-1296.

Woodward, J. (2006). Sensitive and insensitive causation. *The Philosophical Review*, *115*(1), 1-50.

Woodward, J. (2010). Causation in biology: stability, specificity, and the choice of levels of explanation. *Biology & Philosophy*, *25*(3), 287-318.

Woodward, J. (2021). *Causation with a human face: Normative theory and descriptive psychology*. Oxford University Press.