COGNITIVE SCIENCE
A Multidisciplinary Journal

Cognitive Science 45 (2021) e13040

ISSN: 1551-6709 online
DOI: 10.1111/cogs.13040

# The Formal Structure of Kind Representations

Paul Haward,[a] Susan Carey,[a] Sandeep Prasada[b]
[a]*Department of Psychology, Harvard University*
[b]*Department of Psychology, Hunter College*

Received 15 May 2020; received in revised form 31 July 2021; accepted 3 August 2021

**Abstract**

Kind representations, concepts like *table, triangle, dog, and planet*, underlie generic language. Here, we investigate the *formal structure* of kind representations—the structure that distinguishes kind representations from other types of representations. The present studies confirm that participants distinguish generic-supporting properties of individuals (e.g., this watch is made of steel) and accidental properties (e.g., this watch is on the nightstand). Furthermore, work dating back to Aristotle establishes that *only some* generic-supporting properties bear a *principled connection* to the kind, that is, are true of an individual by virtue of its being a member of a specific kind (e.g., telling time for a watch). The present studies tested the hypothesis that principled connections are part of the formal structure of kind representations. Specifically, they tested whether they structure a newly learned kind representation. Experiment 1 found that introducing a property of a newly encountered novel kind in any one of four linguistic frames that provide evidence that a property has a principled connection to a kind (e.g., "It has fur because it is a blick") led participants to infer a *different* conceptual consequence of principled connections (i.e., "There is something wrong with this blick, which does not have fur") for which they had no direct evidence. Two introduction frames that provided no evidence for principled connections (e.g., "Almost all blicks have fur") did not generate the same consequence. Experiment 2 found that all of the targeted properties were generic licensing, irrespective of the introduction frame. That the distinction between properties that bear principled connections to their kinds, and merely generic-supporting properties structures *novel* kind representations, provides strong evidence that this distinction is part of the formal structure of kind representations.

*Keywords:* Concepts; Conceptual representation; Formal structure; Generic knowledge; Kind representations; Modes of explanation

---

Correspondence should be sent to Paul Haward, Department of Psychology, University of California, Berkeley, 2121 Berkeley Way, Berkeley, CA 94704, USA. E-mail: haward@berkeley.edu

## 1. Introduction

Human beings create a vast repertoire of representations for *kinds of things*—concepts like *hammer, bird, library, mountain*, and *universe*. Any normally developing child will generate thousands of kind representations during early conceptual development, often from only a single encounter with a novel entity and without explicit instruction (Bloom, 2000; Markson, Diesendruck, & Bloom, 2008). Once generated, each kind representation sits at a critical interface. On one side of this interface, all kind representations provide a meaningful organization of properties—properties like shape and parts (e.g., legs, a tail, a handle) material constitution (e.g., wood, flesh, blood), and characteristic location (e.g., in a nest, in a factory, in a classroom). For example, for natural kinds, properties are organized according to the principles of psychological essentialism, including representations of causal schema that relate properties of a given kind to one another and explain how new members of the kind come into existence and come to have their kind characteristic properties (Gelman, 2003). On the other side of this interface, kind representations map to nouns in natural language, and as such they provide the meanings for some of the most basic elements of symbolic, unbounded, linguistically expressible thought. For example, they underlie generic uses of language, as in "Birds fly," "Barns are red," and "Ticks carry Lyme disease."

Given the importance of kind representations to commonsense thought and language, investigating their structure and how they are acquired is a central task of theories of conceptual representation and conceptual development (Carey, 1985; Gelman, 2003; Graham & Kilbreath, 2007; Keil, 1989; Macnamara, 1986). Investigating kind representations requires distinguishing between the specific content of the representation of any one kind of thing (i.e., content-specific to the representations of dogs, tables, and trees) and the abstract structure, which is present in *any and all* kind representations. This latter abstract structure is what makes the representations *kind* representations. We call the abstract structure that makes a representation that *type* of representation (in this case, a *kind* representation) its *formal structure*. This structure enters into the generation of each and every kind representation since it is common to all kind representations and characterizes what it means *to be* a kind representation. Here, we investigate if the formal structure of kind representations distinguishes properties that have a principled connection to the kind from properties that are generic-supporting, relative to the kind, but do not have a principled connection to it. We did so by investigating whether participants distinguish merely generic-supporting properties of kinds from principled-connected ones when they first establish kind representations for newly encountered kinds.

### *1.1. The formal structure of kind representations: Characterizing Kinds and their Instances*

All kind representations include a conceptual distinction between properties understood as true of a kind (generic-supporting properties) and properties that are true of some members of a kind or are present in some members of the kind on particular occasions, but which are not true of the kind in general (accidental/idiosyncratic properties; Cimpian, Brandone, & Gelman, 2010; Gelman, 2003; Hollander, Gelman, & Star, 2002; Leslie & Gelman, 2012). For example, *generic sentences* like "Watches tell time," and "Tires are black" reflect an understanding that the property in question in some way characterizes the kind. *Telling time* and *being black* characterize the kinds *watches* and *tires*, respectively, unlike accidental/idiosyncratic properties such as *being dirty* or *being made in Belgium*, which may be true of individual watches or tires but not understood to characterize those kinds. It is important to note that the designation of a property as accidental/idiosyncratic is *kind-relative*. For example, though dirty is an accidental property of a watch, it is a generic-supporting property of junkyards. The distinction between generic-supporting and idiosyncratic/accidentally related properties emerges early in ontogenesis, reflected in the child's language and thought by at least two and a half years of age (Gelman, 2003). Furthermore, at least one study shows that adults are sensitive to the generic/accidental distinction within properties of instances of newly encountered kinds (Sutherland, Cimpian, Leslie, & Gelman, 2015). The ontogenetically early availability of the distinction and its availability to structure novel kinds suggests that this distinction is part of the formal structure of kind representations.

Other evidence shows that characterizing a kind is not simply a matter of representing what is true of *all* or *most* members of a kind (e.g., Carlson & Pelletier, 1995; Leslie, 2007). For example, generics like "Birds lay eggs," "Ticks carry Lyme disease," and "Sharks kill bathers" are judged as true despite the fact that only a small percentage of members of the kinds in question possess those properties. Even when evaluating sentences about the properties that characterize *novel* kinds, participants can judge generic sentences to be true where only a small percentage of members of the kind possess those properties. For instance, when presented with a single sample of lorches where only 10% of the lorches have purple feathers, if the property is presented as distinctive or dangerous, 55% of participants are prepared to judge the sentence "Lorches have purple feathers" as true (Cimpian et al., 2010). Given that these judgments are provided following a single encounter with a set of novel instances of a kind, they likely derive from the formal structure of kind representations.

Within the class of generic-supporting properties for any given kind representation, we observe a further systematic distinction in how properties can relate to a kind. A rich line of thought, tracing back to Aristotle (Charlton, 1970; Witt, 1989), proposes that a subset of generic-supporting properties are understood as existing in instances of the kind *by virtue of* the kind of thing they are (Aristotle in Charlton, 1970; Cudworth, 1731/1996; Haward, Wagner, Carey, & Prasada, 2018; Moravcsik, 1975; Prasada & Dillingham, 2006; Pustejovsky, 1995). These are properties like *telling time* for a *watch*, *barking* for a *dog*, *designing buildings* for an *architect*, and *having three sides* for a *triangle*. These properties—properties that are said to bear "principled connections" to their kinds—are distinguished from merely generic-supporting properties by a variety of conceptual and linguistic signatures. Take "Barns have roofs" and "Barns are red" as examples. Both express true generics. However, participants accept formal explanations for why a particular barn has a roof ("because it is a barn") more so than for why a particular barn is red; judge that there is something *wrong* with a barn without a roof, whereas there is nothing wrong with a barn that is not red; take a singular sentence "A barn has a roof" to express a generic, while "A barn is red" is taken to refer to a particular barn; accept as true "In virtue of being a barn, this has a roof" or "Part of being a barn is to have a roof" more than the comparable statements about being red. All of these differences remain when the prevalence of the properties within their kinds is controlled. These signatures are licensed for the *same* properties of each kind. Furthermore, these signatures are found for kinds and properties from different domains (e.g., having a roof for a barn, telling time for a watch, barking for a dog, designing buildings for an architect; Aristotle in Charlton, 1970; Prasada & Dillingham, 2006, 2009; Prasada, Khemlani, Leslie, & Glucksberg, 2013). Finally, this distinction is present in the child's representations of known kinds by at least age four (Haward et al., 2018).

These results suggest that the distinction between properties that bear principled connections to their kinds and merely generic-supporting properties is part of the formal structure of kind representations. Let us call this possibility the *part of formal structure* hypothesis. However, because the previous research was conducted with known kinds, these results are also consistent with an alternative hypothesis. Rather than being part of the formal structure of kind representations, the so-called signatures of principled connections may be independently learned generalizations about kinds that are acquired on the basis of knowledge about specific kinds and specific properties. We may call this the *independently learned generalizations* hypothesis.

The primary aim of the present studies is to provide evidence that distinguishes between these hypotheses by investigating whether principled connections structure *novel* kind representations. By using novel kinds, we ensure that any systematic pattern of responses must reflect the formal structure of kind representations rather than specific knowledge participants have acquired about the kind. In particular, we investigate whether when participants are given evidence that a property of a novel kind licenses *one* of the signatures of principled connections (e.g., hear "Part of being a blick is having fur"), they immediately represent that property as having a principled connection to the kind and thus expect a *different* signature of principled connections to be licensed for that property for that kind, without any direct evidence for that second signature. In these studies, the different signature is the normative signature, a judgment that there is something wrong with an instance that lacks the property in question (e.g., A blick without fur).

Three features of the current paradigm make it a strong test of the *part of formal structure* hypothesis. First, we targeted properties that are often construed as generic licensing for animal and artifact kinds, such as parts like legs and material constitution like having fur and being made of wood (and we confirmed that participants did indeed construe these properties as generic licensing in Experiment 2). Second, these generic-licensing properties were introduced in one of six different introduction frames, a between-subjects variable. Four of these introduction frames included information that the property licenses a signature of principled connections (e.g., "Part of being a blick is having fur"—designated here "PC-signature frames"). These four frames, each of which introduced the property with a different signature of principled connections, were contrasted with two control frames that included no such information (and one of these two control frames even provided *additional* linguistic evidence that the property is generic licensing, for example, "Almost all blicks have fur"). Our dependent variable was a test of a quite *different* reflection of principled connections—the judgment that there is something wrong with an instance of a kind that lacks the targeted property. This task structure allowed us to investigate whether construing a property as bearing a principled connection to a kind requires evidence *beyond* what is needed to construe it as generic licensing. Third, we investigated whether this distinction between properties that have a principled connection to the kind and properties that characterize a kind but do not have a principled connection to the kind is present in kind representations from the moment they are formed. Evidence that the distinction is present when establishing representations of novel kinds would provide strong evidence that principled connections are part of the formal structure of kind representations.

Though the main goal of Experiment 1 is to investigate the hypothesis that principled connections are part of the formal structure of kind representations, the experiment was also designed to investigate whether the distinction between generic-supporting properties (e.g., being black for a tire) and accidentally related properties (e.g., having dust on it for a tire) is available for structuring novel kinds and thus is also part of the formal structure of kind representations. Previous evidence that this distinction can structure novel kind representations comes from Sutherland et al. (2015). They showed that participants were more likely to misremember a quantified fact as if it had been presented as generic if it was a potentially generic-licensing fact/property for kinds from that domain (e.g., for an animal kind, "All zorbs eat fruits and vegetables") compared to when it was an accidentally related fact/property of individuals from that domain (e.g., "All stups get mud in their hair"). In Experiment 1, we investigated whether participants distinguish generic-supporting and accidentally related properties using the same dependent measure we used to investigate the hypothesis that principled connections are part of the formal structure of kind representations. Participants should be more likely to assign normative value to the generic-licensing properties than to accidental properties since generic-licensing properties could possibly have a principled connection to a kind in the relevant domain, whereas the accidental properties could not. Properties conceived of as accidental relative to a kind could not possibly bear a principled connection to the kind since they do not even characterize kinds.

## 2. Experiment 1

The primary goal of Experiment 1 was to investigate whether principled connections can structure a newly learned kind representation. Participants were introduced to instances of eight novel kinds and taught about a likely generic-supporting property (*targeted property*) of each particular instance (Fig. 1 shows an example trial, in which the individual is a member of the kind blick and the targeted property is having fur).

Participants were introduced to the targeted property via one of six introduction frames. Four of these included a signature of principled connections (*PC-signature frames*)—the Part of introduction frame ("Having fur is part of being a blick"), the Formal Explanation introduction frame ("It has fur because it is a blick"), a bare plural By Virtue introduction frame ("Blicks, by virtue of being blicks, have fur"), and an indefinite singular By Virtue introduction frame ("A blick, by virtue of being a blick, has fur"). The By Virtue signature was introduced in both the bare plural and indefinite singular form to allow us to investigate whether the consequences of introducing a property with a signature of principled connections are stable

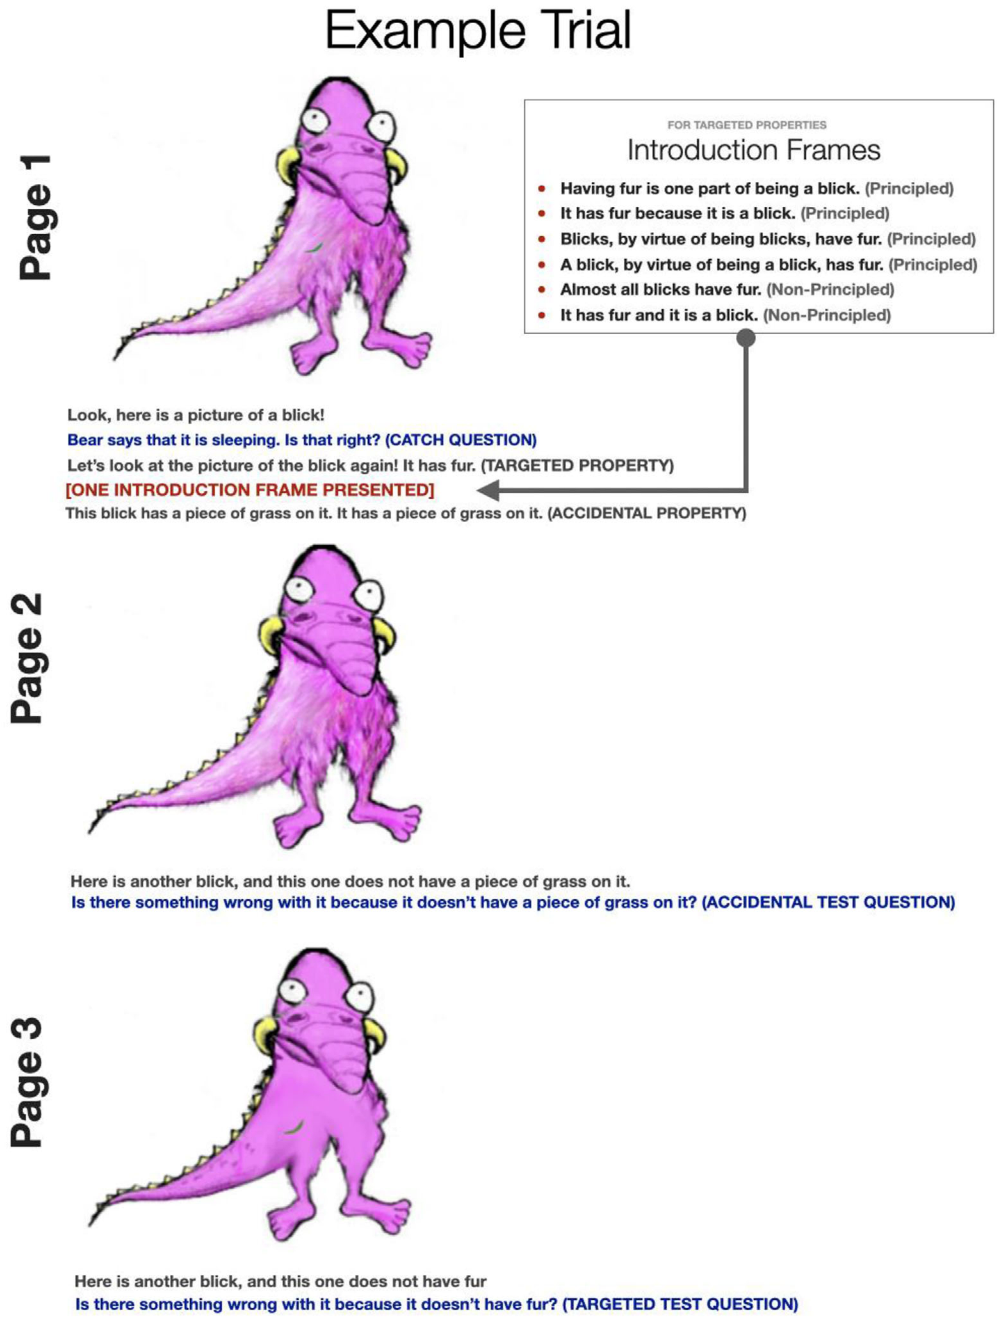


Fig 1. An example trial. Here, the procedure is shown for an animate kind, “blick,” and the property of having fur.

across two different linguistic forms of generics. In addition to the four PC-signature frames, we included two control conditions: a Statistical Generic introduction frame ("Almost all blicks have fur") and a No Information introduction frame ("It has fur, and it is a blick"). The statistical generic introduction frame provides evidence that the property is generic licensing due to its prevalence, but it does not provide specific evidence that the property bears a principled connection to the kind. The No Information frame provides no linguistic information about the relationship between the kind and the property, and certainly no specific information that the property has a principled connection to the kind. However, since all targeted properties were potentially generic-supporting, given the superordinate kinds the novel kinds fell under, we expected that even in the No Information frame the targeted property would be construed as being generic-supporting. On test, participants were introduced to a new instance of the novel kind that *lacked* the targeted property (e.g., a blick without fur). They were then asked whether there was something *wrong* with that instance (e.g., "Is there something wrong with it because it does not have fur").

The *part of formal structure* hypothesis makes the following two predictions. First, the four PC-signature frames should lead to a higher rating that there is something wrong with a member of the novel kind that lacks the targeted property, compared to the two control introduction frames (Statistical Generic and No Information; a preregistered analysis, http://aspredicted.org/blind.php?x=id7un5). Second, participants that are introduced to a targeted property with any of the four PC-signature introduction frames will judge instances that lack the targeted property to have something wrong with them. This is an *exploratory* analysis because ratings on Likert scales can shift around due to exogenous factors (e.g., perhaps participants are hesitant to say there is something *wrong* with these instances across all conditions).

The *independently learned generalizations* hypothesis, in contrast, predicts no effect of the introduction frame on ratings of whether there is something wrong with a kind member that lacks a targeted property. This is because, on this hypothesis, the knowledge that licenses *each* of the signatures associated with principled connections is learned via signature-specific positive evidence for each property and kind, and participants have no relevant evidence for acquiring a normative expectation regarding the targeted property in instances of the *novel* kind.

## 2.1. Method

### 2.1.1. Participants

Five-hundred and ten adults participated online using Amazon's Mechanical Turk. The primary language of all participants was English, and all participants were located in the United States. Participants received a small payment for their participation. A power analysis on pilot data provided a target sample size of 36 participants per condition, which would be sufficient to detect a significant difference ($p < .05$) between the Formal Explanation PC-signature frame and the No Information frame with 80% power. We increased the sample size to 85 participants per introduction frame condition to ensure enough power to detect differences between the introduction frame conditions.[1] We programmed Mechanical Turk to

Table 1
The kinds, targeted properties and accidental properties used for all experiments

| Kind Name | Domain | Targeted Property | Accidental Property |
|---|---|---|---|
| Blick | Animate | Fur | Piece of grass on it |
| Zav | Animate | Scales | 10 years old |
| Wug | Animate | Tail | Eyes open |
| Fep | Animate | Three legs | Next to a rock |
| Jop | Artifact | Made of wood | On a table |
| Mell | Artifact | Rough | Made 4 years ago |
| Liff | Artifact | Handle | Made in New York |
| Timble | Artifact | Three rings | Owned by Bear |

randomly assign each participant to one of the six introduction frame conditions, testing a total of 510 participants. As per our preregistered exclusion criteria, we excluded participants who answered three or more catch questions incorrectly, out of eight total catch questions (described below; this led to exclusion of 12.7% of the participants). Of the remaining participants, 98% of the catch questions were answered correctly. After exclusion, there were 70 participants in the Part of condition, 76 participants in the Formal Explanation condition, 74 participants in the bare plural By Virtue condition, 77 participants in the indefinite singular By Virtue condition, 78 participants in the Statistical Generic condition, and 70 participants in the No Information condition. Of the participants tested, 56% identified as male, 43% identified as female, and 1% preferred not to state their gender.

### *2.1.2. Stimuli and procedures*

Participants viewed eight sets of pictures that depicted instances of four novel animate kinds and four novel artifact kinds (see Fig. 1 for an example of the whole procedure for the animate kind "blick" and the property of having fur). The first picture in each set depicted an instance of the kind and labeled it: "This is a blick." Participants were then asked a catch question, to make sure that the participants were people and were paying attention, in this case, "Is this blick asleep?" Then two properties of the depicted instance were highlighted—a *targeted property* (having fur for the blick) and an *accidental property* (having a piece of grass on it for the blick; see Table 1). The accidental properties were ones that are generally understood to be accidental/idiosyncratic for known animal and artifact kinds. The accidental properties should not be understood as characterizing the kind and thus provide a clear case of a property for which participants should not have any normative expectations that instances of that kind should have that property. The introduction frame for the accidental properties was always the same in all conditions (e.g., "This blick has a piece of grass on it"). In contrast, the targeted properties were chosen because they *might* plausibly license true generics. None was antecedently likely to be an idiosyncratic/accidental property of the pictured instance of the novel kind.

The main independent variable was the frame with which the targeted properties were introduced. See Fig. 1 for examples of the six introduction frames. This was a between-subjects variable and participants were randomly assigned to one of the six introduction frame conditions. We investigated our manipulation of interest (introduction frame) across different domains of kinds (animals and artifacts) and different types of properties (parts and material composition). The hypothesis that principled connections are part of the formal structure of kind representations requires that they should be able to structure different types of novel kinds and different types of properties. Within each domain, two of the targeted properties were *parts* of the objects (e.g., three legs for an animal kind and a handle for an artifact kind) and two were *material* properties (e.g., fur for an animal kind, and being made of wood for an artifact kind). See Table 1 for a full list of the kinds, targeted properties, and accidental properties tested. Including kinds from different domains and different types of properties allowed us to conduct exploratory analyses to see whether participants had any preexisting expectations about which *types* of properties might have principled connections to kinds in each of the two superordinate domains tested (artifacts and animals).

On each trial, after answering the catch question about the first instance of a kind, and having been introduced to the targeted and accidental properties of this instance, participants saw two pictures presented sequentially, each depicting a different instance of the same kind (see Fig. 1). One instance lacked the targeted property and one instance lacked the accidental property. The new instances were introduced as follows: “Look, here is another blick, but this one does not have fur,” or “Look, here is another blick, but this one does not have a piece of grass on it.” After each new instance was presented, the participant was asked to rate on a Likert scale whether there was anything wrong with a blick that did not have that property: For example, “Is there something wrong with it because it does not have a piece of grass on it?” The scale went from −3 (*definitely nothing wrong*) to 3 (*definitely something wrong*). The slider began at 0 (*unsure*), and participants were required to click and move the slider (or confirm its present location) in order to proceed. The presentation order of these two new instances was counterbalanced across trials.

## 2.2. Results

In Experiment 1, the included participants correctly answered the catch trials on 98% of all trials, showing that they were paying attention to the contents of the pictures.[2]

Fig. 2a displays the average ratings that there is something wrong with an individual missing the targeted property as a function of each of the six introduction frames. It also displays the ratings for missing an accidental property, averaged across all six introduction frame conditions, since the introduction frame for the accidental property was identical across all conditions (e.g., “This blick has a piece of grass on it” or “This timble is owned by Bear.”)

### 2.2.1. Distinguishing targeted and accidental properties

As predicted, participants were more likely to judge that there was something wrong with a new instance of a novel kind lacking a targeted property than a new instance of a novel kind lacking an accidental property. Normative violation ratings for instances of kinds lacking an accidental property were significantly lower ($M = -1.96$, $SD = 1.32$) than for instances lacking a targeted property ($M = 0.77$, $SD = 1.48$; $t(444) = 30.10$, $p < .001$). This was the

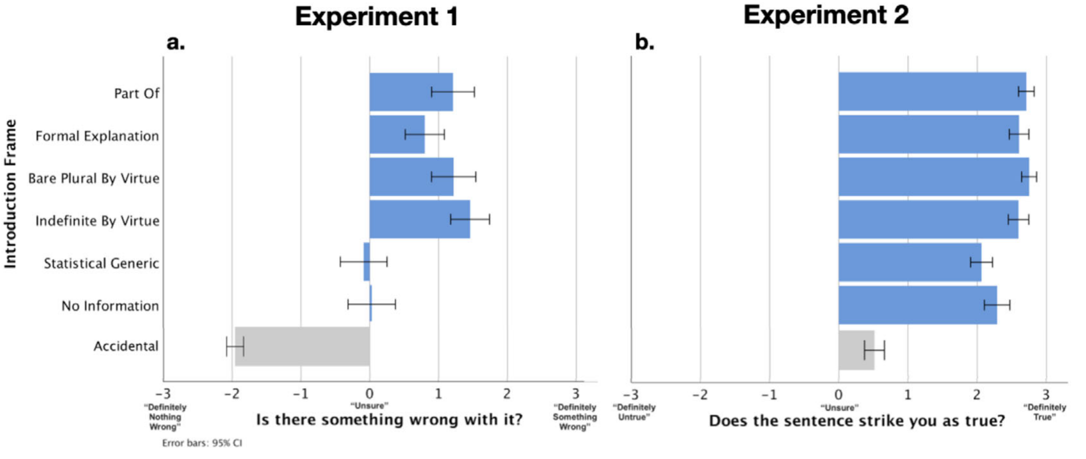


Fig 2. Panel (a) Average normative violation rating for an instance of a novel kind if it lacked a property introduced using each introduction frame. Panel (b) Average truth judgment rating for the generic sentences in Experiment 2, organized by introduction frame condition.

case in every introduction frame condition (*p*s $< .001$). Furthermore, binomial tests showed that the number of participants who judged there to be nothing wrong with instances of kinds lacking an accidental property (had a mean rating of less than 0) was significantly greater than would be expected by chance in every introduction frame condition (all *p*s $< .001$).[3]

### 2.2.2. *Determining the effect of introduction frame*

In order to investigate the role of introduction frame on the normative expectations participants had about the targeted properties, a $6 \times 2 \times 2$ analysis of variance (ANOVA) examined the effects of the introduction frame (Part of, Formal Explanation, indefinite By Virtue, bare plural By Virtue, Statistical Generic, No Information), domain (animate, artifact), and property type (part, material) on the mean normative violation rating for the targeted properties. Our data adhered to the homoscedasticity (homogeneity of variances) assumption underlying parametric tests such as ANOVA but violated the normality assumption. We nevertheless chose to report these statistics as they are widely used and so familiar to a range of audiences. More importantly, they are robust to moderate violations of their assumptions, especially when sample sizes are similar across groups (Glass, Peckham, & Sanders, 1972; Howell, 2009).

There was a significant effect of introduction frame ($F(5,439) = 17.45$, $p < .001$; see Fig. 2a). Post hoc Tukey honestly significant difference (HSD) tests revealed that, as predicted, this was driven by a difference between the four PC-signature frames and the two control frames. Participants were more likely to judge that there was something wrong with an instance of a kind that lacked a targeted property if that targeted property was introduced in any of the PC-signature frames (Part of: $M = 1.21$, $SD = 1.30$; Formal Explanation: $M = 0.80$, $SD = 1.24$; bare plural By Virtue: $M = 1.22$, $SD = 1.39$; indefinite By Virtue: $M = 1.46$, $SD = 1.24$) than if it was introduced in either the Statistical Generic or No Information control frames (Statistical Generic: $M = -.09$, $SD = 1.51$; No Information: $M = .03$, $SD = 1.44$). Furthermore, each of the four PC-signature frames differed significantly from both the Statistical Generic frame and the No Information frame (all $p$s $< .01$). The control frames did not differ from each other ($p = .99$). Finally, the normative ratings from the PC-signature frames did not differ from one another (all $p$s $>$ .4) except for the Formal Explanation frame being rated lower than the indefinite By Virtue frame ($p = .034$).

As expected, participants actively moved the slider away from the point on the scale marked "*unsure*" (0), toward the upper end of the scale in each of the PC-signature frame conditions. The proportion of participants who judged there to be something wrong with instances of a kind (had a mean rating greater than 0) was greater than expected by chance for the Part of condition (0.84), the Formal Explanation condition (0.68), the bare plural By Virtue condition (0.78), and indefinite By Virtue condition (0.86) (all binominal $p$s$<$ .001). Introducing a property with any of the PC-signature frames licenses the inference that there is something wrong with an individual kind member that lacks the property. In contrast, the number of participants who judged that there was something wrong with instances of a kind was not different than expected by chance in the Statistical Generic (proportion = 0.51, $p = .91$) and No Information conditions (proportion = 0.51, $p = 91$). Nevertheless, the mean normative ratings for targeted properties introduced in these frames were differentiated from the normative judgments for instances of kinds lacking accidental properties (e.g., a blick without a piece of grass on it; $p < .001$ in each case).

The inclusion of kinds from different domains (animate and artifact kinds) and different property types (part and material properties) allowed exploratory analyses to investigate whether participants come to this task with expectations about which *types* of properties are likely to have principled connections to kinds for each domain. Analysis of these factors within the main $6 \times 2 \times 2$ ANOVA analysis revealed a main effect of domain ($F(1,439) = 17.81$, $p < .001$), property type ($F(1,439) = 32.98$, $p < .001$) and an interaction between domain and property type ($F(1,439) = 9.99$, $p < .002$). These effects are summarized in Fig. 3. The effect of domain was due to higher normative ratings for animal kinds than artifacts.

This finding adds to previous findings of domain differences in the use of generics for animal and artifact kinds. Both children and adults use more generics when talking about animal kinds than artifact kinds in naturalistic conversation (Gelman, Goetz, Sarnecka, & Flukes, 2008; Goldin-Meadow, Gelman, & Mylander, 2005) and in studies involving novel kinds (Brandone & Gelman, 2009, 2013). The present finding shows that participants understood there to be more principled connections, and thus normative expectations, for animal kinds than for artifact kinds. Importantly though, participants distinguished PC-signature frames from non-PC-signature frames for *both* animal and artifact kinds as predicted by the *part of formal structure* hypothesis. The main effect for property type was due to higher normative ratings for parts than material. The main effects must be interpreted in light of the interactions between these two variables. Normative ratings for instances of a kind lacking a part property did not differ significantly by domain (e.g., artifact kinds: $M = 0.95$; animate kinds: $M = 1.03$; $t(444) = 3.17$, $p = .34$). In contrast, normative ratings for instances of kinds lacking a

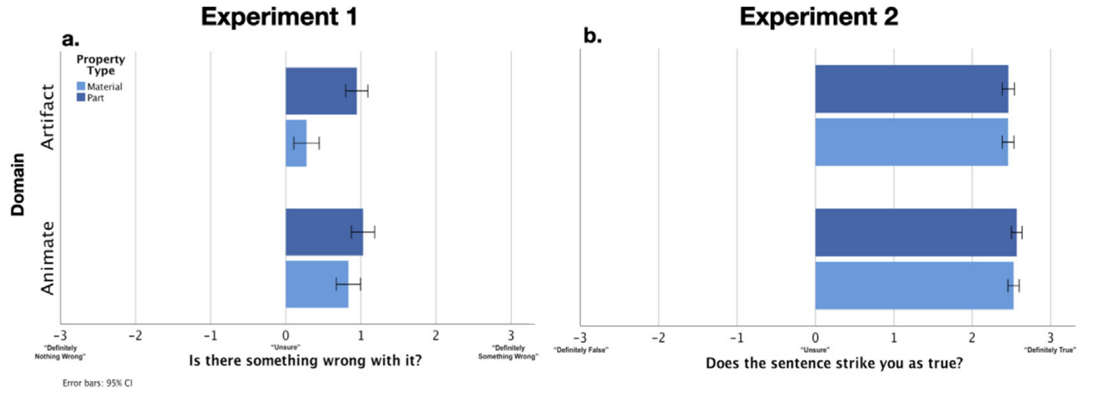


Fig 3. Panel (a) Average normative violation rating for an instance of a novel kind if it lacked a targeted property in Experiment 1, organized by the domain of the novel object and by property type. Panel (b) Average truth judgment rating for the generic sentences in Experiment 2, organized by the domain of the novel object and by property type.

material property were significantly *lower* for artifact kinds (e.g., being made of wood, $M = 0.28$) than animate kinds (e.g., having fur, $M = 0.83$, $t(444) = 7.39$, $p < .001$, corrected for multiple comparisons). These results are consistent with participants bringing knowledge of which types of properties tend to have principled connections to kinds within a domain to the task of learning about novel kinds (Prasada, 2017).

Finally, there was a significant three-way interaction between introduction frame condition, domain, and property type ($F(5,439) = 2.37$, $p = .04$). This was primarily due to the normative ratings in the Formal Explanation condition, which did not differ from the other principled conditions across all trial types, except the artifact material trials, in which the Formal Explanation condition ratings were lower than the other principled conditions. There were no further significant effects.

The *part of formal structure* hypothesis proposes that principled connections structure novel kinds, in general, for different types of kinds and different types of properties. Therefore, we examined whether the PC-introduction frames led to higher normative violation ratings than did the two control frames for each domain type and each property type. The effect of the introduction frame was present across *all* types of trials (see Fig. 4). Four separate one-way ANOVAs, one for each combination of domain and targeted property type (animate material, animate part, artifact material, artifact part) confirmed the main effect of introduction frame on normative ratings for each type of trial (all $p$s $< .001$). Post hoc Tukey HSD tests revealed that in each case, these main effects were due to higher ratings for normative violations in each of the PC-signature introduction frame conditions than in each control frame condition.[4]

## 2.3. Discussion

Experiment 1 provides the first experimental evidence that principled connections can structure *novel* kind representations. Upon hearing the relationship between a property and

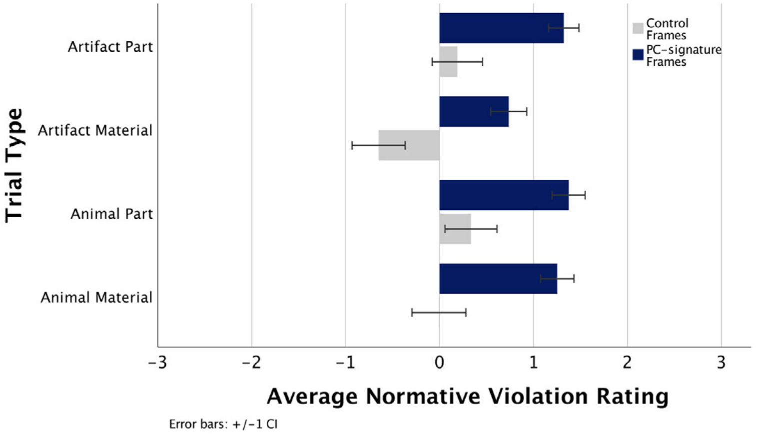


Fig 4. Average normative violation rating for an instance of a novel kind if it lacked a targeted property in Experiment 1, organized by condition (signature of principled connections frame vs. Control frame) and trial type.

a novel kind described with a signature of principled connections, participants form normative expectations regarding that property. Importantly, participants had received no direct evidence concerning whether there is anything wrong with a blick without fur, or a timble without three rings, and indeed, introducing the targeted property in a control frame with no signature of principled connections led to significantly lower normative violation ratings than did PC-signature introduction frames. In the control frame conditions, normative ratings were not differentiated from 0 (*unsure*). Thus, PC-signature introduction frames provided sufficient evidence to establish that the property bears a principled connection to the novel kind. The Part of, By Virtue, and Formal Explanation signatures of principled connections hang together with the normative signature, in spite of the fact that on the surface, the normative signature is very different from the signatures in the PC-signature introduction frames of Experiment 1. Thus, Experiment 1 supports the hypothesis that principled connections between properties and kinds are part of the formal structure that guides the acquisition of each and every new kind representation.

Furthermore, these data provided evidence that the principled connection relation between properties and kinds is distinct from the generic licensing relation between properties and kinds. Properties that were introduced with information that the property/kind relation was generic-supporting in the Statistical Generic introduction frame did *not* subsequently license normative expectations. In fact, the Statistical Generic condition did not differ from the No Information condition, which provided no specific information about the relationship between the kind and the property. Both of these non-principled introduction frame conditions led to participants being at chance levels in judging that there is something wrong with an instance missing targeted properties that are *candidates* for being represented as having a principled connection to the kind given the right evidence. In contrast, and as predicted, instances that are missing properties that are accidentally related to the kind (and thus are not even candidate properties to be represented via a principled connection) were less likely than chance to be judged as having something wrong with them.

Although all the PC-signature introduction frames led to the licensing of normative expectations, introducing a property with the Formal Explanation frame (“It has fur because it is a blick”) did so less than did the other three PC-signature frames. This result, though not predicted, is interpretable. Haward et al. (2018) provided evidence that properties with principled connections to *known* kinds license the Formal Explanation and normative signatures more so than merely generic-licensing properties by age four. But in these studies, for both children and adults, the differentiation between properties that had principled connections to kinds and merely generic-licensing properties was less pronounced for formal explanations than for normative judgments. Haward et al. (2018) speculate that this is because acts of explanation are fundamentally pragmatically driven. People offering explanation seek to provide information they infer the questioner lacks and thus the satisfactoriness of explanations are affected by assumptions about the addressee’s knowledge state and goals in seeking an explanation. Consistent with this, Hemmatian and Sloman (2018) provide evidence that the extent to which formal explanations are judged to answer a question is affected by the extent to which the category label is accepted and used by a community. Recent research by Vasilyeva and Lombrozo (2020) has found that formal explanations (e.g., He is poorly paid because he is an immigrant) are licensed when a category (e.g., immigrants) is understood to occupy a position within society that explains a property (e.g., being poorly paid). It is not currently known if such cases involve principled connections or not, but if they do not, it would be another reason why hearing a formal explanation does not provide as strong evidence for a principled connection between a property and a newly encountered kind than do the other signatures of principled connections.

Finally, Experiment 1 suggests that two sources of evidence are used to determine *which* properties of any new kind representation have a principled connection to the kind. First, the planned analysis of introduction frames suggests that if a property of a novel kind is heard together with a signature of principled connections, then this information can be used to establish the property as having a principled connection to the kind. Second, our exploratory analyses suggested that people may make use of knowledge about which *types* of properties (e.g., material, parts) have principled connections to kinds in different superordinate domains (e.g., animal, artifact) when establishing principled connections between properties and kinds. It is important to note that antecedent knowledge of property/domain relations was not sufficient for establishing a property of a given type as having a principled connection to a basic level kind within that domain. If it had been sufficient, no effect of introduction frame would have been found in the present experiment. Establishing a property as having a principled connection to a novel kind within a domain requires additional evidence. Experiment 1 showed that PC-signature introduction frames can provide that additional evidence.

## 3. Experiment 2

The purpose of Experiment 2 was to test our assumption that all of the targeted properties were taken to be generic licensing within the novel kinds and that the accidental properties were not taken to be so. If this is not so, then we cannot conclude participants distinguish between properties that have principled connections to kinds and merely generic-licensing properties when structuring novel kind representations. Rather, the results of Experiment 1 could be interpreted as PC-introduction frames increasing the likelihood that a targeted property supports a true generic and that this, in turn, increases the normative violation ratings.

Experiment 2 introduced a single individual of eight novel kinds, along with a targeted and an accidental property of each, and a catch question for each, exactly as in Experiment 1. Introduction frame was held constant for each participant. The novel kinds, the properties, and the six introduction frames were identical to those of Experiment 1. After introduction, rather than being shown new instances of the kind that lacked either the targeted or accidental properties, participants were simply asked to rate the truth of generic sentences that stated that the kinds had each property (e.g., "Blicks have fur"). We predicted that ratings in all introduction frame conditions would be high for the targeted properties, robustly above 0 (*unsure*), because parts and material constitution of individuals are in themselves construed as generic licensing properties within animal and artifact kinds. We predicted that the true generic ratings would be much higher for the targeted properties than for the properties we took to be accidental properties of individuals (e.g., "owned by Bear").

### 3.1. Method

#### 3.1.1. Participants

A total of 506 adults participated online using Amazon's Mechanical Turk. The primary language of all participants was English, and all participants were located in the United States. As in Experiment 1, per our preregistered exclusion criteria, we excluded participants who answered three or more catch questions incorrectly, out of eight total catch questions. This yielded 77 participants in the Part of condition, 77 participants in the Formal Explanation condition, 82 participants in the bare plural By Virtue condition, 83 participants in the indefinite By Virtue condition, 80 participants in the Statistical Generic condition, and 77 participants in the No Information condition. Across all participants included in the final sample, 98% of catch questions were answered correctly. Of the participants tested 44% identified as male, 55% identified as female, one person identified with a gender category that was neither male nor female, and 1% preferred not to state their gender.

#### 3.1.2. Stimuli and procedures

The stimuli were identical to those of Experiment 1. The experiment unfolded as in Fig. 1, through page 1. On pages 2 and 3, participants were presented with a bare plural generic sentence about the kind and the targeted property (e.g., "Blicks have fur"), or the kind and the accidental property (e.g., "Blicks have a piece of grass on them"). For each kind-property pair, they were asked: "To what extent does this sentence strike you as true?" Participants provided their responses on a 7-point Likert scale from −3 (*definitely false*) to 3 (*definitely true*) with 0 labeled as "*unsure*." As in Experiment 1, the order of the targeted question and the accidental question was counterbalanced across trials.

### *3.2. Results*

Experiment 2 confirmed that the information in all introduction frames, together with the nature of the targeted properties themselves, establish each kind-targeted property relation as generic licensing. Mean ratings given in each condition are shown in Fig. 2b. Binomial tests confirmed that the number of participants who judged generics like "Blicks have fur" or "Timbles have three rings" to be true (had a mean rating greater than 0) was significantly greater than expected by chance in all the conditions ($p$s $< .001$). The proportion of participants who provided a positive mean rating in each condition was as follows: Part of frame (1.0), Formal Explanation frame (1.0), bare plural By Virtue frame (1.0), indefinite singular By Virtue frame (0.99), Statistical Generic frame (1.0) and No Information frame (0.99). These results contrast sharply with those from Experiment 1. In Experiment 1, it was *only* when targeted properties were introduced in PC-signature introduction frames that the number of participants providing a positive normative rating differed from chance.

Truth ratings for generic statements involving each targeted property were significantly different from those given for generic statements involving the accidental/idiosyncratic properties (e.g., "Blicks have a piece of grass on them" or "Feps are next to a rock": $N = 505$, $M = 0.59$, $SD = 1.60$). This was true in *each* introduction frame condition (all $p$s $< .001$). That this was so in the No Information condition (targeted property generic rating: $M = 2.23$; accidental property generic rating: $M = 0.59$, $t = 9.69$, $p < .001$) shows that the targeted properties themselves were recognized as generic licensing in these kinds.

This differentiation between accidental and targeted properties was also true of the normative judgment dependent variable of Experiment 1. Therefore, across two dependent measures—normative expectations and the licensing of a true generic—adult participants are sensitive to the accidental/generic distinction for a novel kind representation. Nonetheless, and unexpectedly, participants were unsure whether the generic sentences predicating accidental properties were true, with ratings slightly, but significantly, above 0 (see Fig. 2, $M = 0.59$, $t(504) = 8.32$, $p < .001$). A binomial test confirmed that the number of participants who judged generics about accidental properties to be true (provided a mean rating of greater than 0) was greater than expected by chance (prop = 0.62, $p < .001$). Given the kind-relative nature of the distinction between generic-supporting and accidental/idiosyncratic properties, the question arises concerning how we determine whether a given property of an individual should be understood as an accidental/idiosyncratic property of the kind or as a generic-supporting property. Previous research has shown that we can use linguistic cues, nonlinguistic knowledge cues, and pragmatic cues to make generic interpretations (e.g., Cimpian & Markman, 2008; Gelman & Raman, 2003), and thus it is possible that the pragmatics of introducing an instance of a novel kind and pointing out a property may have suggested that the property is lawfully connected to the kind and may be generic-supporting, even though our nonlinguistic knowledge suggests otherwise. More research is needed to understand this small but surprising effect.

An exploratory $6 \times 2 \times 2$ ANOVA investigated the effect of introduction frame (Part of, Formal Explanation, bare plural By Virtue, indefinite By Virtue, Statistical Generic, No Information), domain (animate kind, artifact kind), and property type (part, material) on the average Likert ratings of the truth of the generic sentence. There was a main effect of introduction frame ($F(5,470) = 13.91, p < .001$) and a main effect of domain ($F(1,470) = 10.56, p < .001$). Post hoc Tukey HSD tests revealed that the average generic ratings of the PC-signature frames did not differ from each other (Part of: $M = 2.64$, $SD = 0.56$; Formal Explanation: $M = 2.53$, $SD = 0.70$; bare plural By Virtue: $M = 2.73$, SD = 0.51; Indefinite By Virtue: $M = 2.55$, $SD = 0.73$—all pair-wise $p$s > .05) and neither did the ratings for targeted properties introduced in the two control frames (No Information: $M = 2.24$, $SD = 0.81$; Statistical Generic: $M = 2.06$, $SD = 0.70$, ns). The true generic ratings from each PC-signature frame condition were higher than those from each control frame condition (all $p$s < .05). This is to be expected given that providing evidence that a property bears a principled connection to the kind is strong evidence that the kind/property relation is generic licensing. The crucial feature of these data is that the generic ratings for targeted properties in *all* introduction frame conditions tended overwhelmingly toward the positive end of the Likert scale. This is in contrast to Experiment 1, in which it was only when a targeted property was introduced in a PC-signature frame that the targeted property/kind then licensed the normative signature of principled connections.

Finally, post hoc Tukey HSD tests revealed that the main effect of domain was driven by slightly higher generic ratings for animal kinds ($M = 2.50$, $SD = 0.75$) than for artifact kinds ($M = 2.42$, $SD = 0.80$) consistent with previous findings that children and adults both use more generics when talking about animal kinds than artifact kinds generics (Brandone & Gelman, 2009, 2013; Gelman et al., 2008; Goldin-Meadow et al., 2005).

But in Experiment 2, unlike Experiment 1, there was no hint of an interaction between domain and property type (see Fig. 3). Though exploratory, these results further suggest participants are treating their evaluation of a targeted property's status as generic licensing in Experiment 2 in an entirely different manner from their normative evaluation of the very same targeted properties in Experiment 1.

### 3.3. Conclusions from Experiment 2

Experiment 2 provides evidence that *all* of the introduction frames, together with the nature of the properties themselves, establish the targeted properties as generic licensing. In all introduction frame conditions, including the No Information frame, the mean ratings that the generic statements were true were over 2, where 0 was "*unsure*" and 3 was "*definitely true*." Furthermore, the proportion of participants who judged the sentences to be true was high (between 0.99 and 1.0) and significantly different from chance in all introduction frame conditions. Finally, in all introduction frame conditions, generic ratings for the targeted properties were higher than for the accidental properties, confirming that the distinction between generic-supporting properties and accidental/idiosyncratic properties of a newly encountered individual of a novel kind is available when first learning about that kind (see also Sutherland et al., 2015).

## 4. General discussion

The present studies confirm the conclusion that participants sharply differentiate accidental properties from generic-licensing properties when establishing novel kind representations

(Sutherland et al., 2015). They do so both in their normative judgments ("There is something wrong with a jop not on a table/not made of wood") and their judgments of the truth of bare plural generics ("Jops are on tables/made of wood").

Additionally, the present studies yielded three results that together support the hypothesis that representing principled connections between a kind and some of its properties is *part of the formal structure of kind representations*. First, hearing a property of a novel individual animal or a novel individual artifact in a PC-signature introduction frame led participants to judge there was something *wrong* with a new individual of that kind that did not have that property—a reflection of the *normative* signature of principled connections. This was true despite the fact that no direct evidence was provided for the normative signature. Second, principled connections are *differentiated* from merely generic-supporting properties as soon as a novel kind representation has been generated. Though participants construed all of the targeted properties as generic licensing, when a targeted property was introduced *without* a signature of principled connections, for example, "It has fur and it is a blick," the normative signature was *not* triggered for that property/kind relation. Nor was it triggered when *additional* linguistic generic-licensing information was provided for the relations between the targeted properties and the novel kinds (e.g., "Almost all blicks have fur"). Third, Experiment 1 showed that adults can draw upon at least two different sources of evidence to decide *which* properties may have a principled connection to the kind in a given domain: first, hearing it predicated of the kind with a signature of principled connections, and second, drawing on antecedent knowledge about which types of properties have principled connections to kinds in different domains. Altogether, these results provide strong evidence for the hypothesis that principled connections are part of the formal structure of kind representations; they are available for structuring the representations of novel kinds *when these are first acquired*.

The formal structure of kind representations provides a rich abstract structure for conceiving things as instances of kinds. In the rest of the general discussion, we (i) discuss how the present research deepens our understanding of the nature of principled connections, (ii) identify some questions regarding the acquisition of principled connections raised by the present work, and (iii), discuss other conceptual phenomena that have been attributed to the formal structure of kind representations.

### *4.1. The nature of principled connections between properties and kinds*

The results of the present experiments provide the first definitive evidence that the different signatures of principled connections have an intrinsic unity and are not independently learned generalizations over specific kinds and specific properties. Previous research on principled connections has all involved known kinds and properties. This research could, and did, show that the explanatory, normative, and descriptive dimensions of principled connections co-occur but are left open whether they were acquired together or piecemeal. In demonstrating that these signatures are interlocked, the present experiments provide strong support for an intuition that traces back to Aristotle—that these disparate signatures all reflect a single aspect of kind representations (Ross, 1924).

The unity of the signatures of principled connections that were provided to participants (e.g., Part of, Formal Explanation, By Virtue) and the normative signature that was inferred is remarkable. The former can all be represented in a purely descriptive/explanatory (non-normative) vocabulary, whereas the normative signature clearly cannot. It is intrinsically normative. Part of relations and explanatory relations are typically assumed to be represented in purely non-normative terms. For example, spatial statements involving "part of" are not assumed to make reference to normative terms or have normative force. Neither do formal explanations. How then are the normative and non-normative united in principled connections?

One possibility, first suggested by Aristotle (Anagnostopoulos, 2009; Ross, 1924), is that a type of part-whole understanding underlies this relation between the property and its kind, and as a consequence, these properties are subject to a principle of perfection or completeness; only when an instance of a kind has all of its principledly connected properties is it understood as “complete” (for relevant discussion of these ideas, see also Helmig, 2012; Scaltsas, Charles, & Gill, 1994).[5] What follows is a form of normativity that is not grounded in notions of benefit or moral goodness. We can think of and appreciate the perfect crime or perfect villain just as easily as the perfect gift (for these and other examples see Ross, 1924, Book V). Prasada and Dillingham (2009) provided some empirical support for this proposal, and the present experiments provide further support, since they demonstrate a causal relationship between the Part of signature (e.g., “Having fur is one part of being a blick”) and the normative signature (“A blick without fur has something wrong with it”).

Insofar as principled connections are part of the formal structure of kind representations, this form of normativity is relevant to the representation of concepts for all kinds of things. This suggests that though we can distinguish purely descriptive/explanatory representations from representations that have a normative dimension, the descriptive/explanatory and normative are intrinsically intertwined in kind representations. Other research suggests a similar intertwining of the descriptive and normative in common sense judgments about what is possible (Phillips & Cushman, 2017) and even causal judgments (Knobe, 2010), though the normativity involved in such cases differs from the normativity deriving from principled connections between kinds and properties. It will be important for future research to investigate the ways in which the normative dimension of principled connections may be related to other forms of normativity found in common sense conception (e.g., Knobe, 2010; Knobe, Prasada, & Newman, 2013; Phillips & Cushman, 2017; Roberts, Gelman, & Ho, 2017).

### *4.2. The identification of properties that bear principled connections to kinds*

The relation of bearing a principled connection is kind relative; the property of being red bears a principled connection to the kind strawberry and is merely generic licensing for the kind barn. That principled connections structure *new* kind representations raises a crucial question: How do learners know which properties of a kind should be represented as having a principled connection to a kind?

The studies in the present paper illuminate partial answers to this question, and suggest avenues for future research. Experiment 1 showed that hearing a signature of principled connections is sufficient for marking a property as having a principled connection to the kind. However, further work would be needed to explore whether this mechanism is part of the process by which people *actually* learn which properties have principled connections for a given kind in everyday circumstances. A first step would be to conduct corpus analyses probing the frequency of the signatures of principled connections in speech.

Experiment 1 also suggested that knowledge about which *types* of properties (e.g., material, parts) have principled connections to kinds for kinds from different domains (e.g., animal, artifact) may contribute to whether a person construes a property as having a principled connection to a *novel* kind. Overhypotheses—hypotheses such as *every kind of animal has a characteristic body covering* have been suggested as providing a mechanism through which domain-based knowledge may be used to reason about novel kinds (Goodman, 1955; Kemp, Perfors, & Tenenbaum, 2007; Macario, Shipley, & Billman, 1990; Shipley, 1993). Such overhypotheses could certainly contribute to establishing that a particular property has a principled connection to a novel kind. However, overhypotheses cannot be the sole source of evidence relevant to acquiring principled connections. For overhypotheses to guide the acquisition of principled connections they must be formulated as overhypotheses about properties that have principled connections to known kinds within a domain, in order to distinguish them from overhypotheses involving merely generic-licensing properties (e.g., every kind of service vehicle has its own distinctive color). Unless the hypotheses over which the overhypotheses are formulated are innate, the learner must have some other way of identifying the principled connections between kinds and properties that form the basis for the learned overhypotheses. Therefore, overhypotheses are unlikely to be the only mechanism through which learners learn which properties have principled connections to kinds within a domain (for relevant discussion, see also Prasada, 2017).

Of course, the same observation applies to the linguistic signatures of principled connections being the sole source of information for deciding that a property has a principled connection to a kind for properties of novel kinds. A given property/kind relation must be recognized as a principled connection before it will be expressed with a linguistic signature of that connection. Clearly, further research is needed to establish exactly which types of information may be used to represent a property as having a principled connection to a kind. Research is ongoing in our laboratories exploring whether domain specific causal information (e.g., design information for artifacts) might be sufficient to mark a property as having a principled connection to a kind, as well as the extent to which constraints exist about the types of properties that might be expected to have principled connections to kinds within different domains.

### *4.3. The identification of properties that are generic-supporting within kinds*

"Barns are red" and "Strawberries are red" are true generics, whereas "Chairs are red" is not. Though it was not the focus of our investigation, given the kind-relative nature of the distinction between generic-supporting and accidental/idiosyncratic properties, an entirely parallel set of questions arise concerning how we determine whether a given property of an individual should be understood as an accidental/idiosyncratic property of the kind or as a generic-supporting property. As was the case in the present investigation, we expect that linguistic as well as nonlinguistic cues play a role in identifying whether a property is generic-supporting or accidental/idiosyncratic for a given kind. Children as young as 2 and 3 years old can use syntactic and pragmatic cues to interpret a property as either generic-supporting or as idiosyncratic/accidental. For example, they interpret questions such as *Do birds fly?* as inquiring about whether flying is a generic-supporting property of the kind bird but the question *Do the birds fly?* as inquiring about an idiosyncratic property about the specific birds in a given context (Gelman & Raman, 2003). Furthermore, they interpret *How many legs do they have?* as inquiring about a kind characterizing property if it is used in reference to a single three-legged dog (answering "four") but as inquiring about an idiosyncratic property when used in reference to multiple three-legged dogs (answering "three"; Gelman & Raman, 2003). It is likely that in addition to using syntactic and pragmatic cues, we can use overhypotheses concerning what kinds of properties are generic-supporting in particular superordinate domains of kinds to infer which properties are likely to be generic-supporting or idiosyncratic/accidental (as in the present Experiment 2). It remains for future investigation to more fully determine what the linguistic and non-linguistic cues are and how they interact in identifying properties as generic-supporting or accidental/idiosyncratic for a given kind.

### 4.4. *Conceptual phenomena attributed to the formal structure of kind representations*

The experiments reported here provide evidence that the distinction between principled-connected properties and generic-supporting properties is part of the formal structure of kind representations. They also add to previous evidence that the distinction between accidental/idiosyncratic properties and generic-supporting properties is part of the formal structure of kind representations. This means that learners come to the task of learning a new kind representation seeking to determine which properties of a given thing are accidental or idiosyncratic, and which properties may be true of not just that instance or some instances of a kind, but which characterize the kind generally, and which among those bear a principled connection to the kind.

The formal structure of kind representations is also important for a range of additional conceptual phenomena beyond these two distinctions. First, it likely plays a key role in psychological essentialism that makes the assumption that "people believe a category has an underlying property (essence) that cannot be observed directly but that causes the observable qualities that category members share" (Gelman, 2004, p. 408). It is unlikely that the essence is understood to explain all the properties that category members share (i.e., all the generic-supporting properties of a kind). Instead, it is likely that it is the properties that instances of a kind have by virtue of being that kind of thing are (i.e., have a principled connection to the kind) the properties that are assumed to be explained by the essence.

Second, once a new kind representation is generated, the formal structure of kind representations (e.g., *dog*, *spanner*, *tree*) enters into characterizing instances of that kind in ways that are expressed through true generics like "Dogs wear collars" and "Dogs have four legs" (Gelman, 2003; Hampton, 2012; Strevens, 2000). In doing so, kind representations also provide the means for thinking thoughts about kinds themselves, such as the thought that *dogs evolved from wolves*, which is distinct from the thought that individual dogs evolved from individual wolves (Carlson & Pelletier, 1995; Gelman, 2003).

Third, the formal structure of kind representations also underlies our assumption that there is no limit to the number of instances we can think of as belonging to a kind and that some instances of a kind exist actually, whereas others exist potentially (Macnamara, 1982, 1986; Prasada & Hall, 2019).

Many researchers have suggested that perception provides us with representations of particular things that are not specified for any kind and that these particulars can be individuated and identified in spatio-temporal terms (e.g., Kahneman, Treisman, & Gibbs, 1992; Spelke, 1990; Xu & Carey, 1996). In contrast, thinking and reasoning about those particular things as instances of kinds (e.g., as a table, as some wood) affects the way we individuate and identify instances of that kind across times and situations that does *not* depend simply on spatio-temporal information but depends on the kind of thing we think it to be (e.g., Gupta, 1980; Macnamara, 1986; Wiggins, 1980; Xu, 1997; Xu & Carey, 1996, see also Prasada, 2021; Rips, Blok, & Newman, 2006). As such, a fourth conceptual consequence of the formal structure of kind representations is how it enters into computations of numerical identity (sameness in the sense of same one) of individual entities.

## 5. Concluding remark

As the foregoing makes clear, when we identify an individual as a member of a kind, we do much more than simply group perceptually distinct items together. The formal structure of kind representations provides a perspective for thinking and talking about instances of kinds in which we distinguish accidental/idiosyncratic properties from generic-supporting properties, and merely generic-supporting properties from properties that bear a principled connection to the kind. These distinctions structure kind representations as soon as a new kind representation has been generated. Identifying the formal structure of each of the mental representations possessed by the human mind is a foundational task of the cognitive sciences. By providing a characterization of the formal structure of kind representations, the present research helps clarify the contribution kind representations make to the way humans represent information and the perspectives through which we conceive of the external world.

## Notes

1 The power analysis was conducted on this particular contrast since Haward, Wagner, Carey, and Prasada (2018) provided evidence that the Formal Explanation signature is a weaker measure of principled connections than some of the other signatures. This ensured we had sufficient power to detect differences between each of the four PC-signature conditions and the No Information condition.

2 Original data are available on OSF (https://osf.io/npys3/).
3 The proportion of participants who had a negative mean normative rating was 0.91 in the No Information condition ($p < .001$), 0.88 in the Statistical Generic condition ($p < .001$), 0.84 in the Part of condition ($p < .001$), 0.86 in the Formal Explanation condition ($p < .001$), 0.91 in the bare plural By Virtue condition ($p < .001$), and 0.84 in the indefinite By Virtue condition ($p < .001$).
4 In each of the four ANOVAs, the two control conditions did not differ from each other and Part of, bare plural By Virtue and indefinite By Virtue PC-signature frames led to normative ratings that were significantly greater that each of the two control frames (all *p*s < .01). The Formal Explanation PC-signature frame was less differentiated from the control frames than the other three PC-signature frames for some of the trial types. For the animal material trials the Formal Explanation frame normative ratings were greater than those for both control frames (*p*s < .05). For animal part trials the Formal Explanation frame ratings were greater than those for the No Information frame ($p < .01$) but undifferentiated from ratings in the Statistical Generic frame. For the artifact material and artifact part trials, the Formal Explanation frame ratings were greater than ratings for the Statistical Generic frame (*p*s < .01) but undifferentiated from ratings in the No Information frame. The ratings in the Formal Explanation condition were robustly above the midpoint of the Likert scale in all trial types apart from the artifact material trials as already noted.
5 Aristotle's theory of the part/whole structure intrinsic to kinds, together with his principle of perfection/completeness, is usually (though not always) interpreted as a theory of kinds in the external world (for review, see Anagnostopoulos, 2009). Here, we are explicitly concerned with kind *representations*.

## References


Anagnostopoulos, G. (2009). *A companion to Aristotle*. Chicester, UK: Blackwell Companions to Philosophy.

Bloom, P. (2000). *How children learn the meanings of words*. Cambridge: The MIT Press.

Brandone, A. C., & Gelman, S. A. (2009). Differences in preschoolers' and adults' use of generics about animals and artifacts: A window onto a conceptual divide. *Cognition*, *110*, 1–22.

Brandone, A. C., & Gelman, S. A. (2013). Generic language use reveals domain differences in young children's expectations about animal and artifact categories. *Cognitive Development*, *28*, 63–75.

Carey, S. (1985). *Conceptual change in childhood*. Cambridge, MA: MIT Press.

Carlson, G. N., & Pelletier, F. J. (1995). *The generic book*. Chicago: The University of Chicago Press.

Charlton, W. (1970). *Aristotle's physics: Books 1 and II*. Oxford, England: Oxford University Press.

Cimpian, A., & Markman, E. M. (2008). Preschool children's use of cues to generic meaning. *Cognition*, *107*(1), 19–53.

Cimpian, A., Brandone, A. C., & Gelman, S. A. (2010). Generic statements require little evidence for acceptance but have powerful implications. *Cognitive Science*, *34*(8), 1452–1482

Cudworth, R. (1996). *A treatise concerning eternal and immutable morality*. Cambridge: Cambridge University Press. (Original work published 1731)

Gelman, S. A., & Raman, L. (2003). Preschool children use linguistic form class and pragmatic cues to interpret generics. *Child Development*, *74*(1), 308–325.

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

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

Gelman, S. A., Goetz, P. J., Sarnecka, B. W., & Flukes, J. (2008). Generic language in parent–child conversations. *Language Learning and Development*, *4*, 1–31.

Glass, G. V., Peckham, P. D., & Sanders, J. R. (1972). Consequences of failure to meet assumptions underlying the fixed effects analyses of variance and covariance. *Review of Educational Research*, *42*(3), 237–288.

Goldin-Meadow, S., Gelman, S. A., & Mylander, C. (2005). Expressing generic concepts with and without a language model. *Cognition*, *96*, 109–126.

Graham, S. A., & Kilbreath, C. S. (2007). It's a sign of the kind: Gestures and words guide infants' inductive inferences. *Developmental Psychology*, *43*(5), 1111–1123.

Goodman, N. (1955). *Fact, fiction, and forecast.*. Cambridge, MA: Harvard University Press.

Gupta, A. (1980). *The logic of common nouns: An essay in quantified modal logic*. New Haven: Yale University Press.

Hampton, J. A. (2012). Generics as reflecting conceptual knowledge. *Special Issue of Recherches Linguistiques de Vincennes*, *41*, 9–24.

Haward, P., Wagner, L., Carey, S., & Prasada, S. (2018). The development of principled connections and kind representations. *Cognition*, *176*, 255–268.

Helmig, C. (2012). *Forms and concepts: Concept formation in the platonic tradition*. Berlin-Boston: De Gruyter.

Hemmatian, B., & Sloman, S. A. (2018). Community appeal: Explanation without information. *Journal of Experimental Psychology: General*, *147*(11), 1677–1712.

Hollander, M. A., Gelman, S. A., & Star, J. (2002). Children's interpretation of generic noun phrases. *Developmental Psychology*, *38*(6), 883–894.

Howell, D. C. (2009). *Statistical methods for psychology* (7th ed.). Belmont, CA: Wadsworth.

Kahneman, D., Treisman, A., & Gibbs, B. (1992). The reviewing of object-files: Object specific integration of information. *Cognitive Psychology*, *24*, 175–219.

Keil, F. C. (1989). *Concepts, kinds, and conceptual development*. Cambridge, MA: MIT Press.

Kemp, C., Perfors, A., & Tenenbaum, J. B. (2007). Learning over hypotheses with hierarchical Bayesian models Learning over hypotheses with hierarchical Bayesian models. *Developmental Science*, *10*, 307–321.

Knobe, J. (2010). Person as scientist, person as moralist. *Behavioral and Brain Sciences*, *33*, 315–329.

Knobe, J., Prasada, S., & Newman, G. E. (2013). Dual character concepts and the normative dimension of conceptual representation. *Cognition*, *127*(2), 242–257.

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

Leslie, S. J., & Gelman, S. A. (2012). Quantified statements are recalled as generics: Evidence from preschool children and adults. *Cognitive Psychology*, *64*(3), 186–214.

Macario, J. F., Shipley, E. F., & Billman, D. O. (1990). Induction from a single instance: Formation of a novel category Induction from a single instance: Formation of a novel category. *Journal of Experimental Child Psychology*, *50*, 179–199.

Macnamara, J. (1982). *Names for things: A study of human learning*. Cambridge, MA: MIT Press.

Macnamara, J. (1986). *A border dispute: The place of logic in psychology*. Cambridge, MA: MIT Press.

Markson, L., Diesendruck, G., & Bloom, P. (2008). The shape of thought. *Developmental Science*, *11*, 204–208.

Moravcsik, J. M. E. (1975). Aitia as generative factor in Aristotle's philosophy. *Dialogue*, *14*, 622–638.

Phillips, J., & Cushman, F. (2017). Morality constrains the default representation of what is possible. *Proceedings of the National Academy of Sciences of the United States of America*, *114*, 4649–4654.

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

Prasada, S., & Dillingham, E. M. (2009). Representation of principled connections: A window onto the formal aspect of common sense conception. *Cognitive Science*, *33*, 401–448.

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

Prasada, S. (2017). The scope of formal explanation. *Psychonomic Bulletin & Review*, *24*, 1478–1487.
Prasada, S., & Hall, D. G. (2019). Instance-of-object-kind representations. *Cognition*, *189*, 209–220.
Prasada, S. (2021). *The physical basis of conceptual representation–An addendum to Gallistel* (2020). Cognition, *214*, 104751.
Pustejovsky, J. (1995). *The generative lexicon*. Cambridge, MA: MIT Press.
Rips, L. J., Blok, S., & Newman, G. (2006). Tracing the identity of objects. *Psychological Review*, *113*(1), 1–30.
Roberts, S. O., Gelman, S. A., & Ho, A. K. (2017). So it is, so it shall be: Group regularities license children's prescriptive judgments. *Cognitive Science*, *41*, 576–600.
Ross, W. D. (1924). *Aristotle's metaphysics*. Oxford, England: Clarendon Press.
Scaltsas, T., Charles, D., & Gill, M. L. (1994). *Unity, identity, and explanation in Aristotle's metaphysics*. Oxford, England: Oxford University Press.
Shipley, E. F. (1993). *Categories, hierarchies, and induction*. In D. L. Medin (Ed.), *The psychology of learning and motivation. Advances in research and theory* (Vol. *30*, pp. 265–301). San Diego, CA: Academic Press.
Spelke, E. S. (1990). Principles of object perception. *Cognitive Science*, *14*, 29–56.
Strevens, M. (2000). The essentialist aspect of naive theories. *Cognition*, *74*, 149–175.
Sutherland, S. L., Cimpian, A., Leslie, S. J., & Gelman, S. A. (2015). Memory errors reveal a bias to spontaneously generalize to categories. *Cognitive Science*, *39*(5), 1021–1046.
Vasilyeva, N., & Lombrozo, T. (2020). Structural thinking about social categories: Evidence from formal explanation, generics, and generalization. *Cognition*, *204*, 104383.
Wiggins, D. (1980). *Sameness and substance*. Cambridge, MA: Harvard Univiversity Press.
Witt (1989). *Substance and essence in Aristotle: An interpretation of metaphysics VII–IX*. Ithaca, NY: Cornell University Press.
Xu, F. (1997). From lot's wife to a pillar of salt: Evidence that *physical object* is a sortal concept. *Mind & Language*, *12*(3-4), 365–392.
Xu, F., & Carey, S. (1996). Infants' metaphysics: The case of numerical identity. *Cognitive Psychology*, *30*(2), 111–153.