Cognition 176 (2018) 255–268

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Original Articles

The development of principled connections and kind representations

Paul Haward[a,*], Laura Wagner[b], Susan Carey[a], Sandeep Prasada[c]

[a] *Harvard University, United States*
[b] *Ohio State University, United States*
[c] *Hunter College, City University of New York, United States*

ARTICLE INFO

Keywords:
Conceptual representation
Kind representations
Modes of explanation
Generic knowledge
Principled connections

ABSTRACT

Kind representations draw an important distinction between properties that are understood as existing in instances of a kind by virtue of their being the kind of thing they are and properties that are not understood in this manner. For example, the property of barking for the kind dog is understood as being had by dogs by virtue of the fact that they are dogs. These properties are said to have a *principled connection* to the kind. In contrast, the property of wearing a collar is not understood as existing in instances by virtue of their being dogs, despite the fact that a large percentage of dogs wear collars. Such properties are said to have a *statistical connection* to the kind. Two experiments tested two signatures of principled connections in 4–7 year olds and adults: (i) that principled connections license normative expectations (e.g., we judge there to be something *wrong* with a dog that does not bark), and (ii) that principled connections license formal explanations which explain the existence of a property by reference to the kind (e.g., that barks *because it is a dog*). Experiment 1 showed that both the children and adults have normative expectations for properties that have a principled connection to a kind, but not those that have a mere statistical connection to a kind. Experiment 2 showed that both children and adults are more likely to provide a formal explanation when explaining the existence of properties with a principled connection to a kind than properties with statistical connections to their kinds. Both experiments showed no effect of age (over ages 4, 7, and adulthood) on the extent to which participants differentiated principled and statistical connections. We discuss the implications of the results for theories of conceptual representation and for the structure of explanation.

1. Introduction

Human beings are alone in the animal kingdom in developing an extraordinary repertoire of intricate kind representations—representations for kinds of entities like *dogs*, *watches*, *cities*, *triangles* and *atoms*. Kind representations play a central role in human thought. They underlie the meanings of most count and mass nouns in natural language, and as such, they provide an important interface between non-linguistic conceptual structure and combinatorial, hierarchical, unbounded linguistically expressible thought.

Given the centrality of kind representations in common sense thought and language, investigating the characteristics of kind representations and how they are acquired is a central task in theories of conceptual representation and of conceptual development (Cimpian, 2016; Gelman, 2003; Macnamara, 1986; Margolis, 1998; Prasada, 2016; Xu, 2005, 2012). To investigate kind representations, we must distinguish between the specific content of the representation of any one kind of thing (i.e., information specific to dogs, tables, and trees) and the abstract structure of kind representations which underlies kind representations of any and all kinds of things. The specific content of kind representations, unlike the abstract structure of kind representations, varies from kind to kind. So, for example, the kind representation for dogs will include information that characterizes dogs (e.g., that they are animals, that they bark, have fur, have four legs) and distinguishes dogs from other kinds of things, and the kind representation for tables will include information that characterizes tables (e.g., that they are made by humans, are furniture, have tops, are for putting things on) and distinguishes tables from other kinds of things, and so on for each specific kind of thing. Despite these differences, there is evidence for a common abstract structure underlying the representations of kind concepts in general. This structure is what makes the representations *kind* representations and explicating it provides an abstract characterization of how humans think and speak about any and all kinds of things.

* Corresponding author at: Department of Psychology, Harvard University, 33 Kirkland St, Cambridge, MA 02138, United States.
E-mail address: haward@fas.harvard.edu (P. Haward).

https://doi.org/10.1016/j.cognition.2018.02.001
Received 27 February 2017; Received in revised form 9 January 2018; Accepted 3 February 2018
Available online 31 March 2018

1.1. Characteristics of adult kind representations and evidence that young children's kind representations have at least some of the same characteristics

Most fundamentally, all kind representations have the dual function of providing the means for thinking and talking about entities as instances of kinds, and the means for thinking and talking about kinds, themselves (Prasada, 2016). For example, not only can we form the thought that Fido is an instance of the kind dog, we can entertain thoughts in terms of kinds themselves, such as the thought that dogs evolved from wolves (Carlson, 1980; Gelman, 2003). This latter thought is not the thought that individual dogs evolved from individual wolves, but that dogs as a kind evolved from wolves as a kind. Both of these functions of kind representations are evident early in development. Children's extensions of nouns across multiple individuals, as well as the category-based inductions they make of properties across these individuals, reveal their ability to think of distinct things as being instances of the same kind of thing. Furthermore, by age two and a half, and possibly earlier, children can understand and use generic sentences to think and talk about kinds and the properties that characterize them (e.g., "Birds fly," Gelman, 2003).

Importantly, kind representations allow us to characterize kinds in ways that do not reduce to noting what is true of all, most, or many members of the kind. By at least three years of age, generic statements like *dogs have four legs* and *watches tell time* are understood as attributing properties to kinds and are interpreted distinctly from statements about quantified sets of individuals (e.g., *some/all/most dogs wear collars*) (Brandone, Gelman, & Hedglen, 2014; Hollander, Gelman, & Star, 2002; Leslie & Gelman, 2012). This capacity to predicate properties of kinds is one characteristic of the human endowment for representing kinds. A consequence of the profound difference between kind representations and quantified representations of sets of individuals is an asymmetry in the statistical inferences that *follow* from *learning* a generic generalization and the statistical evidence that supports *judging the truth* of a generic generalization. For example, when participants are introduced to a property of a novel kind using a generic statement (e.g., "Lorches have purple feathers."), they expect the property to apply to nearly *all* members of the kind (Cimpian, Brandone, & Gelman, 2010). Conversely, if they simply learn of particular Lorches that some of them have purple feathers (e.g., 30% or 50% of the Lorches they encounter have purple feathers), they subsequently judge the same generic ("Lorches have purple feathers.") as true, and they do this for a wide variety of percentages (even in conditions where as few as 10% of Lorches have purple feathers). This striking asymmetry in how generic statements are interpreted is a reliable part of the way 4-to 7-year-olds, as well as adults, interpret generics and incorporate property information into their kind representations (Brandone et al., 2014). Thus, these expectations likely reflect the abstract structure of kind representations.

In addition to the distinction detailed above between kind representations and representations of prevalent features of sets of individuals, further aspects of the abstract structure of kind representations distinguish kind representations within specific domains (e.g., the animal kind dog) from representations of quantified sets of individuals (e.g., all existent dogs, all of the dogs who ever played Lassie on the TV series, Lassie). For instance, natural kind representations, including representations of animals, plants, and substances, are structured by the assumptions of psychological essentialism (Gelman, 2003)—the assumption that causally deep, perhaps unknown, features of the members of the kind explain how new members come into existence and explain why those members have their kind relevant properties. These schemata, too, are abstract, early developing, and are made available when we think about things from the perspective of a natural kind.

1.2. A further abstract characteristic of adult kind representations: Principled connections between kinds and properties

Recent work has confirmed Aristotle's observation (Charlton, 1970) that kind representations draw an important distinction between properties that are understood as existing in instances of a kind *by virtue of* their being the kind of thing they are and properties that are not understood in this manner.[1] For example, the properties of *barking* or *having four legs* for the kind dog are both understood as being had by dogs by virtue of the fact that they are dogs: these are among the properties that are understood as making the kind what it is. In contrast, the property of *wearing a collar* is not understood as existing in instances by virtue of their being dogs, despite the fact that a large percentage of dogs wear collars (Prasada & Dillingham, 2006). This distinction is central to kind representations and generalizes to all kinds: for example, for the artifact kind *watch* the property of *telling time* is understood to be true of individual watches because they are the kinds of things they are, whereas the property of having a round face is not understood as being true of individual watches because they are the kinds of things they are, even though we assent to the generic proposition *watches have round faces*.

Properties such as *barking* for *dogs*, or *telling time* for *watches* are said to bear a *principled connection* to the kind (Prasada & Dillingham, 2006, 2009). They are those properties that capture part of what it means to be a member of that kind of thing. And, critically, *all* kinds across *all* domains possess properties that bear a principled connection to the kind. For example, for abstract mathematical kinds, such as triangles, the property of having three sides is understood to have a principled connection to being a triangle, and for social kinds such as architects, the property of designing buildings is understood to have a principled connection to being an architect. Furthermore, principled connections can be distinguished from *statistical connections,* which represent properties that are merely highly statistically correlated with particular kinds. In this context, statistical connections between properties and kinds serve as a control in the search for signatures of principled connections: both support generic generalizations (we assent to "dogs bark" and "dogs wear collars," "watches tell time" and "watches have round faces", and so on, in spite of only the first generic in each pair expressing a principled connection).

Principled connections between properties and kinds have a number of unique conceptual and linguistic consequences. First, properties that bear a principled connection to the kind license *normative expectations* concerning the presence of properties in instances of the kind (Prasada & Dillingham, 2006, 2009). For example, adults judge that there is something *wrong* with a dog that does not bark, but they do not have equivalent expectations for statistical connections: there is nothing wrong with a dog that does not wear a collar. That is, if an instance of a kind lacks a principled property, it is judged to be incomplete or to have something wrong with it, whereas if it lacks a property merely statistically associated with the kind, no such judgment is licensed.

In addition, principled connections license *formal explanations*—references to the kind in order to explain the existence of a property in instances of that kind. For example, we can explain the existence of a property that has a principled connection to a kind in an instance of the kind by simply citing the category: e.g., that thing tells time *because it is a watch.* In contrast, strong statistical relations are not enough; although barns are typically red (a statistical property of barns), explanations that seek to explain a barn's redness by citing the fact that it is a barn were rated as being significantly less natural (Prasada & Dillingham, 2006, 2009; Prasada, Khemlani, Leslie, & Glucksberg, 2013). As with the licensing of *by virtue of* statements and normative expectations, this signature of principled connections generalizes to all domains. For example, we can explain why a person designs buildings by citing the fact that she is an architect.

Though previous research has not explicitly investigated the development of principled connections and these signatures, there are

[1] The distinction in Aristotle regards kinds, not kind representations, but it is kind representations that concern us as psychologists.

some findings which suggest that school aged children may produce formal explanations and have normative expectations for some properties and thus may represent principled connections. For example, research by Taylor, Rhodes, and Gelman (2009) suggests that the formal mode of explanation is a natural mode of explanation for children. Taylor et al. found that 5- and 10-year-old children offered formal explanations to explain the predicted persistence of physical traits in both animals (e.g. "because it is a cow") and human genders (e.g. "because she is a girl") more often than they offered statistical/association based explanations (e.g. "because cows usually have straight tails; because a lot of girls sew"). Five year olds also offered more formal explanations than statistical/association based explanations to explain the persistence of behavior traits associated with animals and human gender whereas the 10-year-olds did not. These data are important because they show that children as young as 5-years-old are able to use formal and statistical modes of explanation. As this research did not seek to distinguish between properties that have a principled connection to a kind from those that have a statistical connection to a kind, it remains an open empirical question as to when children distinguish principled and statistical connections to kinds and when they begin to preferentially provide formal explanations for properties that have a principled connection to the kind as opposed to those that have only a statistical connection to the kind.

Recent research by Roberts, Gelman, and Ho (2017) suggests that children interpret some properties of kinds in a normative manner. Roberts et al. found that 4-year-olds readily interpret properties of social groups as carrying normative force—they judge it not to be okay for members of a group not to conform. Furthermore, they also sometimes use formal explanations (e.g., "because it is a Hibble") to explain why it is not okay for a member of a social group not to conform. These findings make it plausible that 4 year olds may represent properties as having principled connections to kinds. It is unclear, however, if the convergence of normative expectations and formal explanations in Roberts et al. (2017) reveals principled connections as this research was not designed to investigate principled connections and to distinguish them from statistical connections. Furthermore, it is likely that the form of normativity that is relevant in the Roberts et al. (2017) study differs from the normativity involving principled connections. The adults in the Roberts et al. study did not interpret group regularities as carrying normative force and did not disapprove of the nonconformity of a group member; however, adults routinely display normative expectations for properties that have principled connections to kinds (Prasada & Dillingham, 2006, 2009; Prasada et al., 2013).

In addition to licensing normative expectations and formal explanations, principled connections have a range of other conceptual and linguistic consequences (for review, see Prasada & Dillingham, 2006, 2009). The distinction between principled and statistical connections is a formal distinction, in the sense of being a conceptual and/or semantic parameter that is specified for generic licensing properties of kinds, and plays a role in our conceptual representations of kinds that is akin to roles of semantic parameters such as agent vs. patient, or syntactic parameters such as count vs. mass noun in our representations of language. The setting of that abstract parameter (e.g., as a principled connection) then determines the attested signatures of principled connections—licensing "in virtue" statements, licensing the judgment having property x is "part of" or an "aspect of" being in the kind, licensing normative expectations, and licensing formal explanations. These phenomena suggest that the parameter should be thought of as conceptual. But that it might also be a semantic parameter is suggested by the fact that the distinction places constraints on the manner in which generic facts may be expressed within natural language. For example, indefinite singular sentences involving properties with principled connections, such as "A dog has four legs," may be interpreted as a generic statement concerning the kind, like "Dogs bark," whereas such sentences involving properties with statistical connections, such as "A barn is red," are interpreted as a statement about a particular barn, rather than as the generic "Barns are red" (Prasada & Dillingham, 2009).

So long as the child can assign the abstract parameter of being a principled connection relative to a particular property relative to a particular kind, that property/kind relation inherits all of the above consequences. These formal aspects of the child's conceptual system are therefore likely to be of particular importance early in development, when she lacks much domain specific knowledge of relevant causal mechanisms and when she has limited access to a large sample of instances from which to form detailed causal and statistical knowledge, both of which are necessary for detailed understanding of the relations between the kind, on the one hand, and the properties of its instances, on the other.

1.3. *Current studies*

Even though there is extensive evidence that adults establish principled connections between some properties of individuals and those individuals' kinds (Prasada & Dillingham, 2006, 2009; Prasada et al., 2013), thus far there has been no developmental work concerning whether young children do so. The primary aim of the current set of studies, therefore, is to determine whether children, like adults, represent principled connections between kinds and some of their properties, and the degree to which these representations are stable throughout development. In order to test this, we sought evidence for two signatures of principled connections in four-to-seven year old children. Experiment 1 asked whether for 4- to 7-year-old children, like adults, principled connections license *normative expectations*. Experiment 2 asked whether for 4- to 7-year-old children, like adults, principled connections license *formal explanations* to explain the presence of properties that have principled connections to the kind in instances of kinds. In both cases, properties with generic supporting statistical connections were included as a contrast to properties with principled connections to kinds.

These ages were chosen for two reasons. The first is the recent evidence that children as young as age 4 and 5 use kind based formal explanations and make kind based normative judgments (Roberts et al., 2017; Taylor et al., 2009). Additionally, the four-to-seven year old age range in children, along with the data from adults, corresponds to the ages for which other abstract characteristics of our kind-based knowledge have been found to be stable across development (see the above literature review). Therefore, if principled connections are part of the basic machinery humans use to think and speak about kinds, then we should observe continuity between preschool participants through to adults with respect to the two signatures of the abstract formal structure of kind representations probed here. Alternatively, these relatively abstruse signatures of kind representations may be abstracted later in development, awaiting much more experience with the language of kind representations.

In addition to shedding light on the development of kind representations, the present research will inform us of children's ability to make use of the formal mode of explanation. There is abundant evidence that from an early age children make use of causal and teleological modes of explanation (e.g., Callanan & Oakes, 1992; Keil, 1995; Kelemen, 1999; Kelemen, Callanan, Casler, & Pérez-Granados, 2005). But efficient cause and teleological cause are only two of the Aristotelian typology of four causal/explanatory schemata (Charlton, 1970; Moravcsik, 1974, 1975, 1991), and thus there is a need to expand the investigation of whether and when children use the formal mode of explanation which is afforded by kind representations. Furthermore, the present research will expand the investigation of when children systematically incorporate normative expectations into their conceptions of kinds of things.

2. Experiment 1

Prasada and Dillingham (2009) demonstrated that adults are more likely to judge there to be something wrong with an instance of a kind if it lacks a property that has a principled connection to the kind (e.g., there is something wrong with a watch that does not tell time) than if it lacks one that has a statistical connection to the kind (e.g., there is nothing wrong with a barn that is not red), even when the properties in each case are equally prevalent in their kinds. This experiment adapts the methods from Prasada and Dillingham's (2009) Experiment 5 to investigate whether 4- to 7-year-old children show a similar pattern of normative judgments to that of adults tested with the same materials, or, alternatively, whether there are developmental changes in sensitivity to this formal aspect of kind representations over the years 4 to adulthood.

2.1. Method

2.1.1. Participants

Forty-three 4- and 5-year-old children (20 girls and 23 boys; $M = 4$ years 11 months, SD = 7 months), thirty-three 6- and 7-year-old children (23 girls and 10 boys; $M = 7$ years, SD = 5 months), and 36 adults participated. The primary language of all participants was English, although some children spoke additional languages. The families in all of the studies reported here reflected the demographics of volunteers in our database, largely middle class with a stay-at-home parent. Ethnicity for both experiments reported here was largely non-Hispanic American Caucasians (approximately 70%), with Hispanic participants (9%) and other race non-Hispanic participants (21%: Native Americans, Asian, Native Hawaiian and African Americans) making up the rest of the sample. Each child was given a toy or shirt for participating and the parents were given a five-dollar travel reimbursement. The child participants were tested at a university lab. Adults were either recruited as volunteers from the psychology department and tested in the university lab, or recruited and tested as volunteers at a local park.

2.1.2. Stimuli and procedures

Children were introduced to a puppet who joined the child in viewing pictures and who served as the pretend source of questions about them. The pictures were photographs of various kinds of objects, and were presented in a book format with a single picture per page. Each child received eight trials, each trial consisting of two successive pages: an initial page on which an instance of a kind was shown with a prominently displayed property that had a principled or statistical connection to the kind (e.g., a picture of a cow with four legs, a picture of a round plate, respectively) followed by a second page in which the same property was missing (e.g., a picture of a cow with three legs, a picture of a square plate). While viewing the first page, children were asked two questions: a *catch question*, which asked about a manifestly false feature of the picture and tested whether the child was maintaining attention to the task (e.g., for an upright cow, "Look! Here is a picture of a cow, and Bear says that it's lying down. Is that right?"); and a *prevalence question*, which asked children to assess a statement about how common (or prevalent) the target feature was and tested the child's statistical knowledge of the critical properties (e.g., "Bear says that most cows have four legs. Is that right?"). On the second page of each trial, children were shown a picture of the same kind with the principled or statistical connection missing or changed (e.g., a cow with only three legs; a square plate) and they were asked the *normative question* (e.g. "Here is another cow, but this one does not have four legs. So, what is this? That's right! It's a cow that doesn't have four legs. *Is there something wrong with it?*").

By random assignment, half of the children were assigned to a condition in which all of the critical properties had a principled connection to the kind, and half of the children were assigned to a condition in which all the critical properties were statistically connected to their kinds. All trials were of the same type for a given child but within each condition, the order of the eight kinds was counterbalanced. Sessions were videotaped and children's responses were coded offline. The procedure for adults was identical to that used by children, except that the adults' responses were simply recorded by the experimenter.

Most statistical and principled connections were drawn from the Prasada and Dillingham battery of kinds and properties (2006, 2009). These studies yielded six signatures of the distinction between the two generic supporting relations between kinds and their properties, including the findings that principled connections license normative expectations and formal explanations more than do statistical connections—the signatures under study here.

A few individual items were changed to be more appropriate for children. In these cases, the properties were chosen as were those in Prasada and Dillingham's work, initially by intuition: whether they warrant the statements of the form "An x, by virtue of being an x, has property p" (e.g., "a cucumber, by virtue of being a cucumber, is green" vs. "a barn, by virtue of being a barn, is red.") The data from adults in the present studies will serve as a validation of the property-kind contrasts chosen. Appendix A provides the complete list of kind-property pairings in each condition,[2] as well as a list of the catch questions. It also includes an example of the complete experimental script.

2.2. Results

Participants correctly answered the catch questions on 99% of all trials, showing that they were paying attention to the contents of the pictures. To determine if children believed the critical properties were highly frequent for kinds in both conditions, a 3 × 2 ANOVA examined the effects of age group (4- and 5-year-olds, 6- and 7-year-olds, adults) and connection type (principled vs. statistical) on the percentage of prevalence statements accepted across the eight trials. There was a main effect of age ($F(2, 108) = 6.88$, $p < .05$, $\eta_p^2 = .12$); post hoc Tukey HSD tests showed that both groups of children were more likely to agree that most of the instances of the kind had the probed properties than the adults (4- and 5-year-olds, 94%; 6- and 7-year-olds, 98%; adults, 89%). Critically, there was no main effect of connection type and no interaction between age and connection type. That is, both children and adults judged the statistical properties as being just as prevalent for the target kinds (95% and 86%, respectively) as principled properties were (97% and 91%, respectively).

2.2.1. Distinguishing principled and statistical connections

Fig. 1 shows the mean proportion of trials on which an instance of a kind was judged to have something wrong with it if it lacked either a principled or a statistical connection.

A 3 × 2 ANOVA, with participants as the random variable, examined the effects of age group (4- to 5-year olds, 6- to 7-year olds, adults) and connection type (principled vs. statistical) on the proportion of normative violation judgments. There was a main effect of connection type ($F(1, 109) = 129.74$, $p < .001$, $\eta_p^2 = .56$): Participants were significantly more likely to judge there was something wrong with a kind lacking a principled connection (e.g., a cow without four legs, or a watch that didn't tell time; 68% overall) than with a kind lacking a statistical connection (e.g., a cat that didn't scratch furniture, or a

[2] The items involving principled connections in the present study as well as Prasada and Dillingham (2006, 2009) had properties that are true of instances of the kinds generally (e.g. Dogs have four legs). Principled connections need not involve properties that are possessed by the majority of instances of a kind (e.g. Ducks lay eggs) (see Prasada et al., 2013). For principled connections of this latter type, the normative expectations are qualified, although still present. There is nothing wrong with a baby duck or a male duck that does not lay eggs, although there is something wrong with a young post-pubescent, female duck that does not lay eggs.

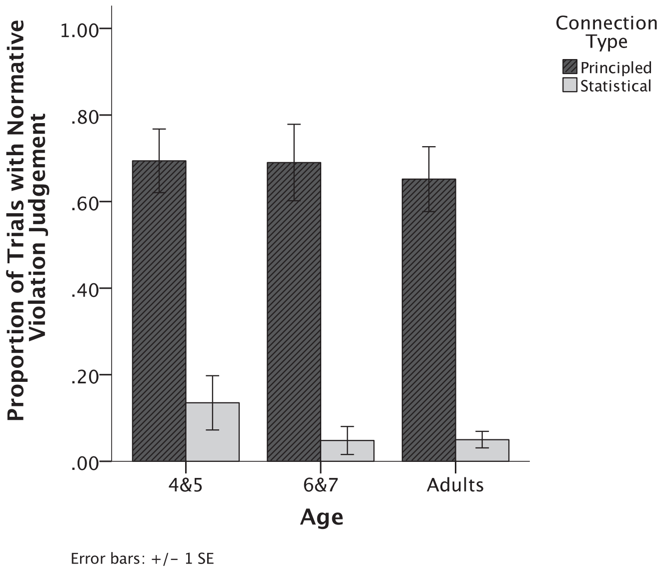

Fig. 1. Mean proportion of trials on which an instance of a kind was judged to have something wrong with it if it lacked either a principled or a statistical property.

school bus that wasn't yellow; 8% overall). There was no main effect of age ($F(2, 108) = .55$, n.s., $\eta_p^2 = .02$) and no interaction between age and condition ($F(2, 108) = .22$, n.s., $\eta_p^2 = .01$). That is, children of both ages differentiated the principled from the statistical properties to the same degree as did adults.[3]

Our measure of participants' beliefs about statistical prevalence of the critical properties is a blunt instrument, namely, whether they assent to the statements "most Xs have property Y." The statement "Most Xs have Y" might be judged as true for any property that is understood as existing in at least 50% of instances of that kind. A consequence of using such a coarse measure of statistical prevalence (chosen because preschoolers might be expected to make sensible judgments) is that it may have masked differences in participants' estimates of the prevalence of the properties with principled- and statistical connections, and these prevalence differences might have driven the difference in normative expectations found in Experiment 1. In order to address this concern, we ran an additional study on MTurk in which we asked 45 adults to provide prevalence estimates (from 0 to 100%) for each of the kind-property pairings used in Experiment 1 (e.g., "What percentage of watches do you think tell time?"). Items with principled connections and statistical connections did indeed differ in the percentage of instances adults understood as having that particular property (91% for properties with principled connections vs 70% for properties with statistical connections; $t(44) = 13.26$, $p < .001$). However, we were able to generate a subset of three properties with principled connections and three properties with statistical connections to their kind that were matched on mean estimated prevalence (84% vs 80%).[4] When the same ANOVA analysis was performed for this subset of items, the results were qualitatively the same as with the full set. Importantly, there was a main effect of connection type ($F(1, 109) = 89.83$, $p < .001$, $\eta_p^2 = .46$): Participants were significantly more likely to judge there was something wrong with a kind lacking a principled property (e.g., a bird that could not fly, or a cucumber that wasn't green; 67% overall) than with a kind lacking a statistical property (e.g., a school bus that wasn't yellow, or a fire truck that wasn't red; 9% overall) even though the items did not differ in the estimated prevalence of their critical properties. There was a marginal main effect of age ($F(2, 108) = 3.07$, $p = .051$, $\eta_p^2 = .06$), which was driven by a lower rate of normative judgments overall for the adults (25%) compared to either the 4- and 5-year-olds (47%) or the 6- and 7-year-olds (41%). Critically, there was no interaction between age and condition ($F(2, 108) = 1.09$, $p = .34.$, $\eta_p^2 = .02$); each age group distinguished principled and statistical connections to the same degree. These results, together with those of Prasada and Dillingham (2006, 2009) which also controlled for prevalence, demonstrate that statistical prevalence cannot account for the difference in the normative expectations assigned to properties that have principled and statistical connections to kinds.

2.3. Discussion

This experiment provided strong evidence for a core signature of principled connections between kinds and properties, namely, that they license normative expectations. Both children and adults were significantly more likely to judge that there was something wrong with an object lacking a property that had a principled connection to the kind than an object lacking a property that had a statistical connection to the kind, even though participants judged both the principled and the statistical properties as being highly prevalent within their kinds. That is, while participants believed that both square plates and three-legged cows are unusual, they only believed there was something wrong with such cows—not such plates. It is the nature of the connection of the property to the kind that matters, not the frequency with which the kind possesses the property. Moreover, unlike the form of normativity investigated in Roberts et al. (2017), this experiment found strong continuity between children and adults in their normative expectations regarding the presence/absence of properties with principled connections to kinds. The differentiation in normative judgments between principled and statistical connections was robustly found in the youngest age group (4- to 5-year-old children) and remained consistently strong for the older children as well as for adults. These results therefore strongly support the idea that an adult-like understanding that only properties that bear a principled connection to the kind license normative judgments is firmly in place by the pre-school years.

Experiment 1 provides evidence that principled connections are part of children's cognitive machinery for thinking and speaking about kinds. In order to gain a fuller grasp on the status of principled connections in the young child's kind representations, we looked at a second signature of principled connections—that principled connections license *formal explanations*. Experiment 2 looked to test for this signature of the distinction between principled and statistical connections between properties and kinds in years 4 to adulthood.

3. Experiment 2

Prasada and Dillingham's (2006, 2009) studies on formal explanation asked participants to rate explanation goodness. Adults were presented with a question, for example, "Why does that [pointing to a dog] have four legs?", and then presented with a number of explanations, one of which was a formal explanation (e.g., *because it's a dog*), and asked to what extent the explanation was "a good or natural response to the question" on a 7-point scale. In Experiment 2, we probed children's (and adults') propensity to provide formal explanations in the simplest possible way—by asking our participants for explanations. Experiment 2 elicited explanations with questions about why individuals had particular properties, some of which should elicit formal explanations because the properties have a principled connection to the kind (e.g., "Why are both of these [pointing to pictures of two knives] sharp?"), and some which should not elicit formal explanations because the properties have merely a statistical connection to the kind (e.g., "why are both of these (pointing to pictures of two plates) round?"). Thus,

[3] A parallel ANOVA, using items as the random variable rather than participants, produced the same results: a main effect of connection type ($F(1, 15) = 176.45$, $p < .001$), no effect of age ($F(2, 15) = 2.305$, n.s.), and no interaction ($F(2, 15) = .11$, n.s.).

[4] The items with principled connections were: birds, cucumbers and knives. The items with statistical connections were: bricks, fire trucks, and school buses.

while Prasada and Dillingham's data provide evidence that adults rate formal explanations for properties that have a principled connection to a kind as being better than formal explanations for properties that have a statistical connection to the kind, the present experiment is the first to explore whether adults actually *provide* more formal explanations for properties that have principled connection to kinds than those that have statistical connections. Moreover, we probed children's explanations in exactly the same direct way as we do adults' without the potential added complications of a switched-at-birth task (cf. Taylor et al., 2009) allowing us to determine whether children also spontaneously provide formal explanations, and whether they distinguish between principled and statistical connections to the same extent as do adults. Alternatively, there may be developmental changes in sensitivity to this formal aspect of kind representations over the years 4 to adulthood.

3.1. Method

3.1.1. Participants

Sixty-two 4- and 5-year-old children (31 girls and 31 boys; $M = 5$ years, SD = 7 months), fifty-three 6- and 7-year-old children (28 girls and 25 boys; $M = 6$ years 11 months, SD = 8 months), and 40 adults participated in Experiment 2.

Approximately half of the child participants in each age group were tested at a university lab in Cambridge, MA. These children were recruited via birth records and through a local children's museum and families were provided with a token incentive (e.g. a small toy) and a small travel reimbursement for participating. The other half of child participants were tested at a lab housed within a science museum in Columbus, OH. These children were recruited from among the museum's visitors and received no incentives for their participation. All adults were tested in MA, and were either recruited as volunteers from the psychology department and tested in the university lab, or recruited and tested at a local park.

3.1.2. Explanation elicitation task

3.1.2.1. Stimuli and procedures. As in Experiment 1, children were introduced to a puppet who joined the child in viewing a set of pictures and who served as the nominal source of all the questions. Stimuli for the experiment were eight sets of color photographs of two instances of a kind (e.g., pictures of two cucumbers, side by side). Each child received eight trials, and each trial contained a different kind of object. In addition, two properties were chosen for each kind of object, a property that does not support generics and so would not be expected to elicit formal explanation (for the control question) and a generic supporting property (for the critical question). The control questions ensured that children were paying attention, demonstrated they could provide relevant causal explanations, and allowed us to verify that they were not perseverating by giving a formal explanation for every question asked. Half of the critical properties had principled connections to their kinds, and half had statistical connections to their kinds. Children were randomly assigned to the principled connection condition (in which all critical properties had principled connections to their kind) or to the statistical connection condition. The kind-property pairings for critical questions were the same as in Experiment 1, and are displayed in Table A2, along with the control properties participants were asked to explain.

On each trial, the experimenter read a short vignette that described a causal event resulting in each object having the same new property (e.g., "These knives are very dirty after chopping vegetables all day. The chef decided to put them in the sink. He put this knife *in the sink*, and he also put this knife *in the sink*"). The vignette then introduced a *critical property*—a property with a principled connection to the kind for children in the principled connection condition (e.g. "This knife is *very sharp*, and this knife is *very sharp*"), or a property with a statistical connection to the kind for children in the statistical condition (e.g., for plates: "This plate is *round* and this plate is *round*").

The vignette was followed by the two questions as to why "these," indicating each pictured object, have each property—the critical property first (e.g., "Why do both of these (watches) tell time?"), followed by the control property (e.g., "Why do both of these (watches) have sand in them?"). The critical question and control question order was fixed for each item, but counterbalanced across items. See Appendix B for the full script of a trial.

Sessions were videotaped and children's answers were transcribed offline for later coding.

Adults were tested with the same materials and procedure as children with two exceptions. Adults were told that the study was designed primarily for young children and were encouraged to give the answers that came most naturally to them, and they were not videotaped. Rather their responses were transcribed by the experimenter.

3.1.2.2. Coding. Explanations were coded into five categories: *non-explanations*, plus four different types of explanation (*formal*, *causal*, *functional*, and *statistical*).

Formal explanations made reference to the kind in order to explain the existence of the property. For example, if a participant responded to the question, "Why are both of these (cucumbers) green?" with the answer, "Because they're cucumbers", their response was coded as a formal explanation. Furthermore, some responses made tacit reference to a norm (e.g., "Because they're *supposed* to be that way") or an abstract category (e.g., "Because they're the same *kind of thing*"). These too were coded as formal explanations.[5]

Causal explanations made reference to a causal process that brought about the property. For example, if a participant responded to the question, "Why are both of these (school buses) yellow?" with the explanation, "Because someone painted them", then the explanation was coded as causal.

Functional explanations made reference to an end to which the property is put. For example, if a participant responded to the question, "Why do both of these (watches) tell time?" with the answer, "So that people can plan their day", then the explanation was coded as functional.

Finally, *statistical explanations* explicitly made reference to statistical facts about the kind using quantifiers such as *all*, *most*, *some* or *always*. For example, if a participant responded to the question, "Why do both of these [watches] tell time?" with the answer, "Because all/most/some watches tell time," then the explanation was coded as statistical.

All other responses that did not directly answer the question, such as uncertain responses (e.g., "don't know"), responses with irrelevant information (e.g., "'cause the frog is" when asked why two plates are round), and tautologies (e.g., "because they tell time," when asked why two watches tell time) were coded as *non-explanations*.

Eighteen percent of the 2480 explanations were coded by 2 coders. Intercoder agreement was 90%. The majority of disagreements involved explanations that were ambiguous between causal explanations and functional explanations (e.g., "because they like it" as an explanation for why two cats scratch furniture). Importantly, agreement was 95% on the designation of formal explanations. Disagreements were resolved by discussion.

3.2. Results

Both adults and children virtually always provided relevant explanations. On the Control questions, 100% of the adults' and 92% of the children's responses to the why question were actually relevant explanations and for the Critical questions 98% of the adults' and 91% of the children's responses were so as well. Fig. 2 displays the distribution of explanation types for control questions and critical

[5] The same pattern of results is found if these latter two types of responses are not coded as formal explanations.

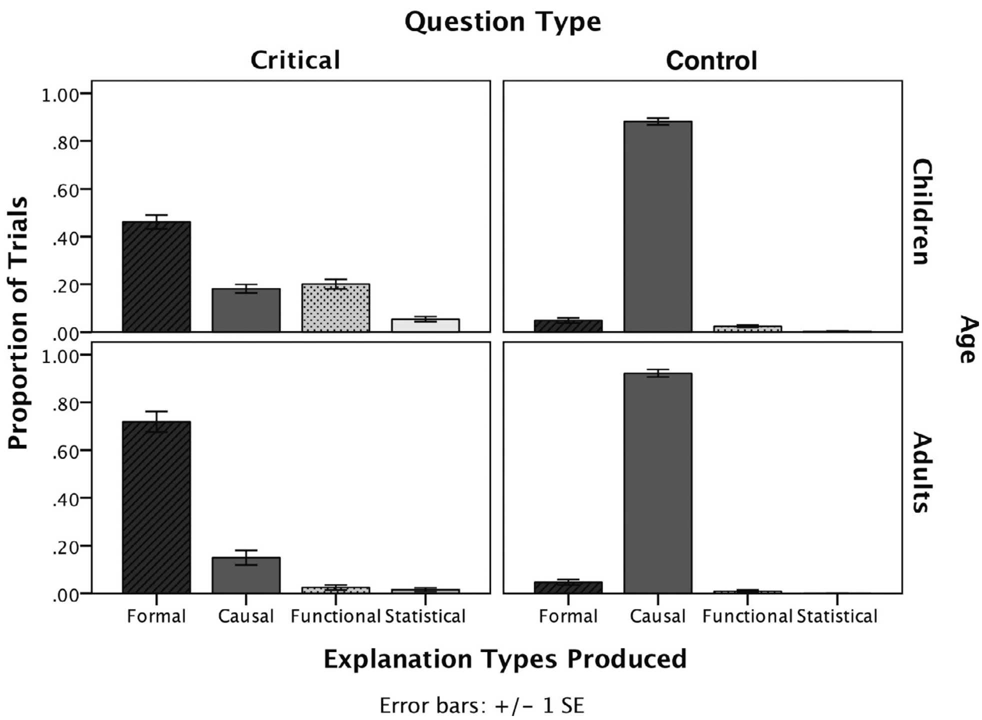

Fig. 2. Distribution of explanation types for control questions and critical questions, collapsing over the Statistical and Principled conditions, and collapsing over age groups among the children.

questions, collapsing over the Statistical and Principled conditions, and collapsing over age groups among the children.

3.2.1. *Distribution of explanation types for critical and control questions*

As can be seen from Fig. 2, the vast majority of the relevant explanations provided on the Control questions were causal explanations in which the participant articulated causally relevant aspects of the process through which the property in the explanandum came to be. Furthermore, the children's explanations contained the same content as the adults'. For example, when provided with a description of two friends at a beach whose watches get blown into the sand and asked of the two watches, why do both of these have sand in them, a typical adult explanation was "because they were blown into the sand by the wind"; a typical 4-year-old explanation was "because they got blown in the sand" and a typical 7-year old explanation was "because the wind blew them off the towel." The responses to the control question establish that the children in this experiment were adept at inferring explanatory relations from information presented in the accompanying vignette, and importantly, they do not simply provide formal explanations in response to every why-question.

Fig. 2 also shows that both children and adults found it natural to explain the existence of the critical properties by citing the kind—i.e., by providing formal explanations. Asked why both of these (cows) have four legs, or are round (plates), participants were quite likely to reply "because they are cows" or "because they are plates." Indeed, the *most* frequent explanation type for the critical questions for both children (46% of relevant explanations) and adults (72% of relevant explanations) was formal explanation. As Fig. 2 shows, for both children and adults critical questions also elicited functional explanations (e.g., for cows, "so they can walk" as an explanation for why both of these have four legs), or causal explanations (e.g., for plates, "because they were made that way in a factory" as an explanation for both of these being round). Statistical explanations were markedly less frequent.

In sum, both children and adults spontaneously produced formal explanations at a high rate for critical questions, and neither did so for the control questions. These data, as well as those of Taylor et al. (2009), confirm that formal explanation, in addition to teleological explanation and to causal explanation that specifies efficient cause, is a *natural* mode of explanation for adults and young children. We now turn to our main question of interest: whether children, like adults, distinguish principled connections between kinds and their properties from merely statistical connections, and offer formal explanations more for properties with principled connections than statistical connections to kinds.

3.2.2. *Distinguishing principled and statistical connections*

Fig. 3 shows the mean percentage of trials on which a formal

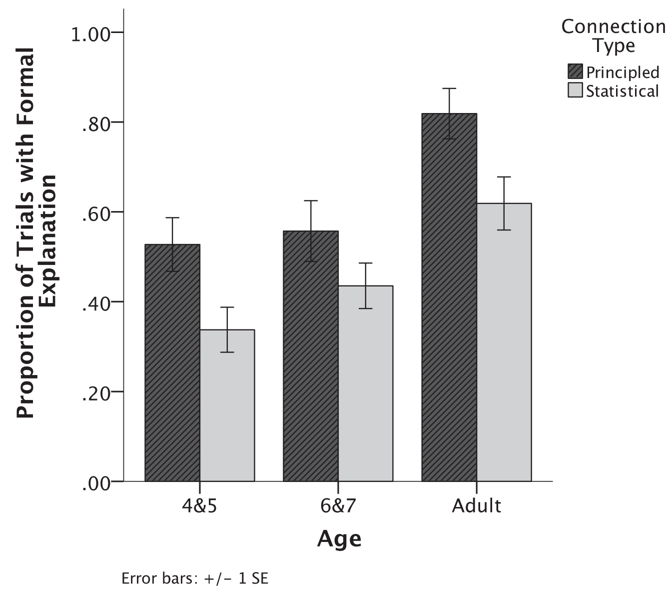

Fig. 3. Mean percentage of trials on which a formal explanation was given for both the principled property condition and for the statistical property condition, broken down by age.

explanation was given for both the principled connection condition and for the statistical connection condition, broken down by age. A 2 × 3 ANOVA examined the effects of connection type (principled, statistical), and age group (4- and 5-year-olds, 6- and 7-year-olds, adults) on the percentage of the eight trials in which a formal explanation was provided. There was a main effect of condition: Participants provided more formal explanations for properties that bear a principled connection to the kind (61%) than properties that bear a statistical connection to the kind (44%: $F(1, 152) = 12.579$, $p = .001$, $\eta_p^2 = .08$). There was also a main effect of age: Adults gave more formal explanations than did children (adults = 72%, children = 46%, $F(2, 151) = 11.988$, $p < .001$, $\eta_p^2 = .14$).[6] Importantly, there was no interaction among these variables; children differentiated principled connections from statistical connections as much as did adults. Furthermore, as can be seen in Fig. 2, there were no developmental changes in the likelihood to provide formal explanations, or in the differentiation of principled and statistical connections, over the ages of 4–5 to 6–7. An ANOVA on the children alone confirmed a main effect for connection type ($F(1, 112) = 7.402$, $p = .008$, $\eta_p^2 = .06$), and no main effect or interactions involving age group.[7]

3.3. Discussion

Experiment 2 confirms the large literature establishing that young children have a grasp of many aspects of the structure of explanation (e.g. Callanan & Oakes, 1992; Gelman, 1990; Keil, 1995; Kelemen, 1999; Kelemen et al., 2005) and are capable of providing formal explanations (Roberts et al., 2017; Taylor et al., 2009). They overwhelmingly produced relevant causal and teleological explanations, with the same content as the adult explanations, for the control questions, and occasionally for the critical questions as well. Experiment 2 yielded two important new results. First, like adults, 4- to 7-year-old children spontaneously produced many formal explanations. Both children and adults find formal explanation natural, but neither group deployed it indiscriminately. Both groups used it only to explain properties that support generic statements, such as “plates are round” or “watches tell time,” but not properties that were in the particular scenarios alone, like “knives are in the sink.” Second, to the same extent as adults, 4- to 7-year-olds distinguished generic generalizations that express principled connections between kinds and their properties from those that expressed merely statistical connections between kinds and their properties, providing more formal explanations for the presence of properties with principled connections to their kinds.

Two further aspects of the data are worth noting: first, neither children nor adults provided formal explanations for principled connections on all trials (children, 54% of trials: adults, 81%), and second, both groups provided formal explanations for statistical connections in a large percentage of trials (children, 39%: adults, 61% of trials).

A lower-than-ceiling rate of formal explanations for principled properties is consistent with the picture of kind representations outlined in Prasada and Dillingham (2006, 2009). Our kind representations allow properties that bear a principled connection to the kind to be explained using a formal mode of explanation. Nevertheless, such a theory of kinds does not *dictate* that a given individual’s having a particular property be explained by that individual’s membership in the kind. Other modes of explanation, such as the teleological mode of explanation and the causal mode of explanation, are available to our participants and may be drawn upon independently to explain the existence of the properties presented in our stimuli. For example, that a property, such as *telling time,* bears a principled connection to the kind *watch* does not prevent the child from drawing upon other explanatory modes in order to explain why a given watch tells time. Children who did not provide a formal explanation for the time-telling of watches instead provided either a functional explanation (e.g., “so they can know what time it is”) or a causal explanation (e.g., “because they’re made to do it”) in order to explain why watches tell time.

What we did not predict was the substantial percentage of trials in which a formal explanation was given to explain a *statistical* property, both by children and adults. Children’s unexpectedly high proportion of formal explanations for the statistical items could be due to the fact that children of these ages may be beginning to represent the distinction between principled and statistical connections between kinds and their properties, and may take any generic supporting property of a kind to license formal explanations, or may mistakenly think that some items that involve statistical connections involve principled connections. This would explain why children produced a higher than expected number of formal explanations for the statistical items. However, the results from Experiment 1 undermine this explanation. Experiment 1 found developmental continuity between age 4 and adulthood in the degree of differentiation of properties that have a principled connection to a kind from properties that have a statistical connection to a kind for the very same properties with respect to normative expectations. Furthermore, this account could not explain why the *adults* produced an even larger number of formal explanations for statistical items. We must explain why *both* children and adults provided so many formal explanations for properties merely statistically connected to their kinds. It is an important task for future research to determine the factors that elicit formal explanations from both children and adults when there is only a statistical connection between the kind and property to be explained. We speculate that several features of the present task elicited formal explanations in both the principled property and statistical property conditions of Experiment 2. Several of these derive from pragmatic factors that influence an explainer’s decision of what information an explanation-seeker lacks.

Given that an individual may be understood to be an instance of many kinds (e.g. Fido is a collie, a dog, an animal), one role of a formal explanation is to identify the kind, if any, to which a property is lawfully connected and thus the connection that requires explanation. This function of formal explanation is relevant to explaining any property that licenses generics, including generics for which the connection between the property and the kind is merely statistical, as in “Barns are red.” A second role of formal explanations is to actually explain the property that has a lawful connection to the kind via a formal connection between the kind and the property—by the property being represented as an aspect of being that kind of thing (Prasada & Dillingham, 2009). This second role is only possible when there is a principled connection between the kind and the property (see General Discussion for a more detailed discussion of the nature of formal explanation). It is possible that participants sometimes produced formal explanations for properties that merely have statistical connections to a kind because in such cases one of the roles of formal explanation is satisfied, namely the role if identifying the relevant kind. This possibility also may have been amplified by pragmatic factors specific to the task we used.

All acts of explanation have a major pragmatic component: the explainer seeks to provide information he or she infers the explanation seeker lacks (e.g., see Salmon, 1999). It is likely that the critical questions in the current task support the inference that the questioner might not know the kind that is relevant to the property in question (whether the connection is principled *or* statistical). The pragmatic context of explanation in these experiments is one in which participants are asked to explain the existence of a property in multiple instances of the same kind (e.g., why are *both* of these yellow? pointing to two bananas). This context may make membership in the same kind salient, because it highlights the fact that the property is true of multiple instances of the

[6] A parallel ANOVA, using items as the random variable rather than participants, produced the same results: a main effect of connection type ($F(1, 28) = 8.35$, $p = .007$), a main effect of age ($F(2, 27) = 17.39$, $p < .001$), and no interaction ($F(2, 27) = .71$, n.s.).

[7] A parallel ANOVA, using items as the random variable rather than participants, produced the same results: a main effect of connection type ($F(1, 28) = 8.92$, $p = .006$), no effect of age ($F(2, 27) = 1.50$, n.s.), and no interaction ($F(2, 27) = .42$, n.s.).

same kind, making it likely that the questioner is asking what kind is relevant to the property (e.g., for school buses and yellow—vehicle (no), bus (no), school bus (yes)). If this is the case, asking participants to explain the presence of a property in a single instance of a kind (e.g., Why is this yellow?) should lead to fewer formal explanations being generated. A second pragmatic factor that may have suggested that the questioner was ignorant of the kind that was relevant to the property was that the question used a pronoun (e.g., Why are both of *these* yellow?) creating a strong pressure to inform the questioner of the kind of thing being referred to. If this is the case, replacing the pronoun with the kind name (e.g., Why are both of these *school buses* yellow?) should lead to fewer formal explanations being generated.

Future research should determine whether these (or other) factors influence the likelihood of formal explanations for all properties that license generic statements, including properties that merely have a statistical connection to a kind. Importantly, despite the factors that make formal explanations natural in explaining why individuals have properties that have *both* statistical and principled connections to their kinds, we found that principled connections provide *greater* warrant for formal explanation than do statistical connections, for 4- to-7-year-olds to the same extent as for adults.

4. General discussion

Experiments 1 and 2 yielded five results that constrain our theories of kind representations, explanation, and conceptual representations more generally. First, Experiment 1 showed that 4- to 7-year-old children, as well as adults, judge there to be something wrong with an instance of a kind lacking a property that has a principled connection to a kind, whereas they deny that there is something wrong with an instance of a kind lacking a property that has a statistical connection to a kind. Second, Experiment 2 showed formal explanation to be a natural mode of explanation, both for young children and adults. Third, Experiment 2 showed that four- to seven-year-old children, as well as adults, are more likely to provide a formal explanation when explaining the existence of properties of two instances of a given kind with a principled connection to that kind than they did for properties with statistical connections to their kinds. Fourth, both experiments showed no effect of age on the measures of differentiation of properties that have principled and statistical connections to kinds: sensitivity to this distinction is stable during development from age 4 through adulthood. Fifth, the differentiation of principled and statistical properties was more categorical in Experiment 1 than in Experiment 2; suggesting that the distinction's implications for normative judgments are greater than its implications for the likelihood of producing formal explanation.

4.1. Nature and development of kind representations

Previous work has identified several features that distinguish kind representations from other types of mental representations (of properties, of events, of logical connectives, of spatial relations, of quantifiers, and so on). These include the capacity for setting up kind representations from only one or two encounters with a novel object (Markson, Diesendruck, & Bloom, 2008), thinking of instances of kinds as being equivalent with respect to indefinitely many other actually as well as potentially existing instances (Prasada, 2016), providing the basis for placing novel objects under the assumptions of psychological essentialism (Gelman, 2003, 2004), and supporting generic generalizations, which reflect our understanding of properties attributed to kinds as opposed to quantified sets of individuals (Gelman, Sánchez Tapia, & Leslie, 2015; Leslie & Gelman, 2012), which are in turn reflected in the asymmetric inferences concerning generic sentences and prevalence expectations discussed in the introduction (Cimpian et al., 2010).

Another distinctive feature of our kind representations is that they mark principled connections to properties that are true of instances of a kind by virtue of their being the kind of thing they are. Previous research with adults has shown that principled connections license formal explanations of properties in instances of kinds and support normative expectations concerning the presence of properties in instances of the kind (Prasada & Dillingham, 2006, 2009; Prasada et al., 2013). The present studies extend this work by showing that adults actually produce formal explanations for principled connections (not merely rate formal explanations as acceptable) and answer categorically that there is something wrong with an instance of a kind lacking a property that has a principled connection to a kind. More importantly, the data from the present experiments show that, like the other central features of kind representations, these signatures of principled connections are in place by age 4, in essentially their adult form. Even the youngest children distinguished principled connections from merely statistical connections to the same degree adults did.

4.2. Implications for the normative dimension of common sense concepts

Experiment 1 showed that both children and adults make a categorical distinction between principled and statistical connections with respect to their engendering normative expectations. Participants virtually never said there was something wrong with an individual that lacked a property merely statistically connected to its kind, whereas they frequently judged that there was something wrong with an individual that lacked a property with a principled connection to its kind. This was equally true from age 4 to adulthood. Thus, Experiment 2 provides the first developmental evidence that the kind representations in place by age 4 specify normative expectations that individuals should have the properties that have principled connections to the kind.

The existence of such normative expectations poses a significant challenge for theories of conceptual representation and development. How does the child, or adult, go from observing what *is* the case in one or more instances of a kind to conclusions about what *is supposed to be* the case for those and other instances of the kind? This problem is similar to the is-ought problem identified by Hume (1738); however, the normativity in the present case is not moral normativity. We speculate that the solution to the present version of Hume's problem is that the normativity derives from the abstract formal structure of kind representations, although further research is needed to determine the nature of the normativity embodied in kind representations. Perhaps, since properties that have a principled connection to a kind are understood to be *aspects* of being that kind of thing, the normativity is grounded in a principle of completeness or perfection (Bublitz & Prasada, 2013; Prasada & Dillingham, 2009).

4.3. Implications for theories of explanation

Explanation plays an important role in human cognition (e.g., Keil & Wilson, 2000; Lombrozo, 2006, 2016). Analyses of the why questions preschool children ask (e.g., Callanan & Oakes, 1992), the answers they give to why questions (e.g., Wellman, Hickling, & Schult, 1997), as well as studies of what explanations young children find satisfying (e.g., Frazier, Gelman, & Wellman, 2016) establish that young children seek, produce, and understand causal and teleological explanations. Furthermore, much research on conceptual representations highlights the central role of explanation in structuring kind concepts, both for adults and young children. Psychological essentialism posits a commitment to their being deep causal properties that explain how new individuals of a kind come into being and explain why they have some of the properties they do (e.g., Gelman, 2003; Strevens, 2000). Experimental work on kind learning shows that participants weigh both the causally central and functionally relevant properties more in their judgments about which other novel entities are members of the kind (e.g., Ahn, Kim, Lassaline, & Dennis, 2000; Lombrozo, 2009; Medin & Shoben, 1988; Sloman, Love, & Ahn, 1998). Dating back to Aristotle, philosophers have noted that there is another mode of explanation, in addition to causal and teleological explanation, namely, formal explanation, which has a central place in mathematical explanation. Indeed, the first modern philosophical theory of explanation, the nomological-deductive theory of Hempel and Oppenheim (1948) sought to reduce all scientific explanation to providing premises whose truth is accepted and showing how the explanandum follows, deductively or inductively, from those premises and accepted laws. Explaining a property of an instance by citing its membership in a relevant kind is an instance of formal explanation that is licensed by the formal (principled) connection between a kind and a property. The data from Experiment 2 confirm Taylor et al.'s (2009) observation that young children, as well as adults, have access to a *formal mode of explanation*, in addition to the *causal* and *teleological* modes, and provide systematic evidence that, as for adults, formal explanation plays an important role in kind representations from age 4.

In fact, our results revealed that formal explanation is a remarkably natural mode of explanation for both children and adults. In Experiment 2, formal explanation was the most common type of explanation drawn upon to explain critical properties (i.e., generic licensing properties such as the property of *telling time* for two instances of the kind *watch* and the property of *being yellow* for two instances of the kind *school bus*). This was true of both children (46% of relevant explanations) and adults (72% of relevant explanations). This was true in spite of the fact that, for each of the critical properties tested, other modes of explanation could have been used to explain the existence of the critical property. And indeed they were: on 20% of trials children provided a functional explanation, as did adults on 3% of trials, and on 18% of trials children provided a causal explanation, as did adults on 15% of trials. For example, in explaining the four-leggedness of a cow (a property with a principled connection to this kind), participants might have referred to either causal origin (e.g., because it's coded in their DNA, because it was born from cows), function (e.g., so that they can walk), or statistical prevalence (e.g., because most cows have four legs) using explanatory and linguistic resources that we know they have access to by age 4 (Gelman, 1990; Kelemen, 1999). Similarly, for yellowness of schoolbuses (a property with a statistical connection to this kind) participants might have appealed to causal origin (e.g., because they were painted yellow), function (e.g., so they can be identified as schoolbuses) or statistical prevalence (e.g., because most schoolbuses are yellow). Despite this, formal explanations were the most common explanation type for both adults and children.

One question raised by these data is why this was so. Formal explanations, unlike the other modes of explanation, appear not to provide additional information about the existence of a property in an object. For example, causal explanations provide information about the process of production of a property, while functional explanations provide information regarding the environmental circumstances and goals of the particular thing (or its designer) that may indicate the reasons for a property's existence. In contrast, formal explanations only provide a reference to the kind, and so at first blush, they appear to be a relatively uninformative mode of explanation.

There are two roles that formal explanation might play that would explain our participants' readiness to provide them in response to a request to explain a property's existence. First, formal explanation may play a role in indicating which kind should be drawn upon to explain the existence of a given property in multiple instances. Most entities fall under a number of different kind descriptions. For example, if a person is encountered in Central Park on a summer's day, then she may be thought of and spoken about using the kind concepts *person*, *mother*, *American*, or *artist*, and the application of each of these kinds on a particular occasion will make a different set of properties intelligible by virtue of the individual being that kind of thing (e.g., for the concept *artist*, the tools she possesses, and for the concept *person*, the number of legs she has). As such, producing a formal explanation may provide information about the most general kind for which the property to be explained is lawfully connected to the kind. This role for formal explanation applies to properties that are merely highly statistically connected to their kinds, as well as applying, of course, to those with principled connections to their kinds. In explaining why this thing is red, it is relevant that it is a barn, not that it is a building. In so doing, a formal explanation conveys the kind, if any, to which a property is lawfully connected and thus the connection that requires explanation. The second role of formal explanations is to actually explain the property via a formal connection between the kind and the property—by representing the property as an aspect of being that kind of thing (Prasada & Dillingham, 2009).

These suggestions make sense of an unexpected feature of our data. Whereas Prasada and Dillingham (2006, 2009) suggest that formal explanation is licensed *only* by principled connections, in Experiment 2, formal explanation was also the *dominant* mode of explanation for accounting for properties of individuals that are merely statistically connected to their kinds, although to a lesser extent than was so for principled connections. That formal explanation may highlight the relevant kinds under which properties have reliable connections to those properties partly explains why they are sometimes natural for properties statistically connected to the highlighted kinds and as well as why they are more natural for properties with principled connections to the kind. This highlighting function may have been amplified by two characteristics of the task used in Experiment 2. First, the questioner did not identify the kind of the instances that had the properties to be explained (e.g., "Why do these have property x?"), putting pragmatic pressure on the explainer to identify the relevant kind. Secondly, participants were not asked to explain why a single instance had a property, but why multiple (two) instances of the same kind had a given property, suggesting a lawful connection between the kind and the property. Nonetheless, the consistent (across ages) differentiation of principled and statistical properties with respect to the likelihood of formal explanations reflects additional explanatory warrant that derives from the formal connection itself—the representation of the property as an aspect of the kind.

4.4. Two crucial open questions

Experiments 1 and 2 suggest that at least by age 4, children's acquisition of kind concepts is constrained by the formal framework provided by kind representations first fleshed out in Prasada and Dillingham (2006, 2009). That is, children and adults alike approach the problem of concept acquisition distinguishing kind representations from other types of concepts. They expect that some of the properties of a given kind will have principled connections to that kind, and that these are properties that exemplars of that kind are supposed to have, and the presence of these properties in instances of a kind can be explained by reference to the kind of thing something is. The research presented here, as well as that in previous work on adults (Prasada & Dillingham, 2006, 2009; Prasada et al., 2013) leaves open two (related) crucial questions. First, for a given kind, which properties have a principled connection to that kind? Though there is work on what types of properties (e.g. entrenched, causally deep) may characterize kinds (e.g. Goodman, 1955/1983; Gelman & Markman, 1986; Shipley, 1993), that research does not address the question of which properties of a kind have a principled connection to that kind. This is because not all properties that characterize a kind involve principled connections to the kind. Statistical connections can license generics (e.g., barns are red) and causal connections (e.g., deer ticks carry Lyme disease) can also license generics, however, neither of these property-kind relations display the characteristics of principled connections. There is nothing wrong with a barn that is not red or a deer tick that does not carry Lyme disease; being red is not a part/aspect of being a barn and carrying Lyme disease is not a part/aspect of being a deer tick. As such, the question of which properties of a kind have a principled connection to that kind remains. The second crucial open question is, what types of information do people use to *establish which* properties have principled connections to which kind. The present study shows that not only do 4-

year-old children distinguish principled connections from statistical ones, they have also identified the principled connections probed in these studies as such. How do 4-year-olds, or adults, for that matter, classify yellow as a statistical property of school buses, and green as a principled property of cucumbers?

Both of these questions are important topics for further research. Concerning the first question, it is likely that there is no single non-formal property that is shared by all principled connections. Nevertheless, it remains possible that there are systematic mappings between different types of non-formal properties and principled connections in different domains (Prasada, 2017). It will be important to develop and explore these possibilities in future work. With respect to the second question, how children and adults *decide* which properties have principled connections to their kinds, it is possible that learners come to the task of concept formation with expectations about the types of properties that kinds of different types (i.e. from different domains) have principled connections to (Prasada, 2017; Prasada & Dillingham, 2009). Another, not mutually exclusive, hypothesis is that the linguistic reflections of principled connections in speech provide relevant evidence. For example, hearing an indefinite, singular generic such as “a dog has 4 legs,” or a formal explanation such as “that has 4 legs because it’s a dog,” or an explicit normative statement such as “dogs are supposed to have four legs” may allow learners to assume that a principled connection is involved and assume the other characteristics of principled connections. Work on the important question of how participants determine which generic supporting properties have principled connections to their kinds is ongoing in our laboratories.

4.5. Conclusion and directions for further research

Human beings, from a young age, are capable of acquiring a vast number of richly structured kind representations during conceptual development, often from limited evidence. This capacity derives in part to the learner’s sensitivity to the different ways in which properties are understood in relation to the kind (Prasada et al., 2013). The two experiments reported here provide evidence that by the preschool years, children represent some properties of individuals as related to their kinds via principled connections, and that sensitivity to the distinction between principled connections and statistical connections is stable throughout development from age 4 through adulthood. Furthermore, Experiment 1 provides evidence for early emerging normative implications of principled connections, and Experiment 2 provides evidence for the importance of formal explanation in the child’s explanatory repertoire, and for its greater naturalness when explaining why individuals have properties with principled connections to their kinds, relative to those with statistical connections to their kinds. The present research also raises a rich set of further questions, including the following: How do children and adults determine which properties of a given kind have principled connections to the kind? Work with adults shows that they distinguish principled connections from both statistical connections and causal connections (Prasada et al., 2013), and the present studies show that children distinguish principled and statistical connections by age 4. When do children *begin* to make these distinctions? What is the exact nature of the normative expectations implicated by kind representations? How does the knowledge structured by our kind representations complement and interact with other types of conceptual knowledge in the course of conceptual development? The answers to these and related questions promise to enrich our understanding of conceptual representation and development.

Acknowledgements

The authors would like to thank Cai Guo and Taylor Wong, who helped with data collection and analysis and Elaine Dillingham who helped with stimulus construction for this paper. This research was supported by a grant from the James S. McDonnell Foundation (#2200204418) and a Harvard University grant from The Foundations of Human Behavior Initiative (FHB), awarded to Susan Carey.

Appendix A. Stimuli and script for Experiment 1

See Table A1.

Table A1
Kinds, critical properties and catch questions for Experiment 1.

Kind	Critical Property	Catch Question: Bear Says…Is Bear right?
Principled condition		
Bird	Flies	Of a bird flying in the air, “it is swimming.”
Cow	Four legs	Of an awake, upright, cow, “it is sleeping.”
Zebra	Black and white stripes	Of zebras, “zebras make honey.”
Cucumber	Green	Of a typically shaped cucumber, “it is square.”
Ambulance	Sirens	Of ambulances, “ambulances go ‘moo’.”
Boat	Floats on water	Of a typically designed boat, “it has wheels.”
Knife	Sharp	Of a black-handled knife, “it is purple.”
Watch	Tells time	Of a black and silver watch, “it is green.”
Statistical condition		
Cat	Scratches furniture	Of cats, “cats bark.”
Dog	Chases sticks	Of a typical dog, “it has wings.”
Mouse	Eats cheese	Of a mouse with its eyes open, “it is asleep.”
Carrot	In bags	Of a bag of orange carrots, “they are green.”
Brick	Rectangular	Of a typical, rectangular brick, “it is soft.”
Fire truck	Red	Of a parked firetruck, “it is flying in the air.”
Plate	Round	Of a white plate, “it is purple.”
School bus	Yellow	Of a parked school bus, “it is driving in water.”

A.1. Script for Experiment 1

The following is the script for a sample trial from the statistical condition in Experiment 1.

[*Experimenter points to the puppet and says:*] This is my friend Bear. Would you like to play a game with Bear? Bear went to the library and he got his book with all of these cool pictures, and he would like to use it to play a game with you. Would you like to play a game with him?

Great let me tell you how Bear plays this game.

In this game he is going to show you some pictures and tell you about them. You have to listen carefully, okay because Bear is going to need your help to figure out lots of different things.

Bear is a bit silly, so sometimes he says things that are right, and sometimes he says things that are wrong, so he's going to need your help to figure out when he says something right, and when he says something wrong.

You think you can help him figure out when he's right and when he's wrong? Great! Are you ready to begin?

Catch Question: [*Experimenter points to the picture of a school bus and says:*] This is a picture of a school bus and Bear says that it's driving in water. Is that right? [*Child provides answer.*]

Prevalence Judgment: This school bus is yellow. Bear says that most school buses are yellow. Is that right? [*Child provides answer and Experimenter turns page*].

Here's a picture of another school bus but this one is not yellow. So, what is this? [*Experimenter corrects child if needed.*] Is there something wrong with it? [*Child provides answer.*]

Appendix B. Stimuli and script for Experiment 2

See Table A2.

B.1. Script for Experiment 2

The following is the script for a sample trial from the principled condition in Experiment 2.

[*Experimenter points to the puppet and says:*] This is my friend Bear. Would you like to play a game with Bear? Bear went to the library and got this book with all these cool pictures, and he would like to use it to play a game with you.

Would you like me to tell you how to play?

In this game, Bear is going to show you some pictures and tell you things about them. You have to listen carefully, because Bear is going to need your help to figure out lots of different things.

Bear doesn't understand some things about the pictures he is going to show you, so he is going to need your help to explain them to him.

Can you help explain some things to Bear? Great! Are you ready to begin?

Table A2
Kinds, critical properties and control properties for Experiment 2.

Kind	Critical Property	Control Property
Principled condition		
Bird	Flies	Hungry
Cow	Four legs	In the barn
Zebra	Black and white stripes	Thirsty
Cucumber	Green	Frozen
Ambulance	Sirens	Wet
Boat	Floats on water	Smells like paint
Knife	Sharp	In the sink
Watch	Tells time	Has sand in it
Statistical condition		
Cat	Scratches furniture	Out of breath
Dog	Chases sticks	Standing
Mouse	Eats cheese	Thirsty
Carrot	In bags	Salty
Brick	Rectangular	Hot
Fire truck	Red	Shiny and clean
Plate	Round	Smells like fish
School bus	Yellow	Dented

[*Experimenter points to the picture of two watches and says:*] Amber and Mary went to the beach. Before they went in the ocean, they took off their watches, and put them on their towels. A big wind came up and blew the watches off the towels and into the sand. So this watch got blown into the sand, and this watch got blown into the sand. And this watch tells time, and this watch also tells time.

[*Critical question*]: This one tells time, and this one tells time. Can you tell Bear why both of these tell time? [*Child provides explanation.*]

[*Control question*]: And this one has sand in it, and this one also has sand in it. Can you tell Bear why both of these have sand in them? [*Child provides explanation.*]

Appendix C. Supplementary material

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.cognition.2018.02.001.

References

Ahn, W., Kim, N. S., Lassaline, M. E., & Dennis, M. J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology, 41*, 361–416.

Brandone, A. C., Gelman, S. A., & Hedglen, J. (2014). Young children's intuitions about the truth conditions and implications of novel generics and quantified statements. *Cognitive Science*.

Bublitz, D., & Prasada, S. (2013). Normative dimensions of generic knowledge. *Poster presented at the 35th annual cognitive science conference, July 31–Aug. 3, Berlin, Germany*.

Callanan, M. A., & Oakes, L. (1992). Preschoolers' questions and parents' explanations: Causal thinking in everyday activity. *Cognitive Development, 7*, 213–233.

Carlson, G. N. (1980). *Reference to kinds in English.* New York: Garland Publishing.

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

Cimpian, A. (2016). The privileged status of category representations in early development. *Child Development Perspectives, 10*(2), 99–104.

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.

Frazier, B. N., Gelman, S. A., & Wellman, H. M. (2016). Young children prefer and remember satisfying explanations. *Journal of Cognition and Development, 17*(5), 718–736.

Gelman, R. (1990). First principles organize attention to and learning about relevant data: Number and the animate-inanimate distinction as examples. *Cognitive Science, 14*, 79–106.

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., & Markman, E. M. (1986). Categories and induction in young children. *Cognition, 23*, 183–209.

Gelman, S. A., Sánchez Tapia, I., & Leslie, S. J. (2015). Memory for generic and quantified sentences in Spanish-speaking children and adults. *Journal of Child Language*, 1–14.

Goodman, N. (1983). *Fact, fiction, and forecast* (4th ed.). Harvard University Press (Original work published in 1955.).

Hempel, C., & Oppenheim, P. (1948). Studies in the logic of explanation. *Philosophy of Science, 15*, 135–175.

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

Hume, D. (1738). *A treatise of human nature*.

Keil, F. C. (1995). The growth of causal understandings of natural kinds: Modes of construal and the emergence of biological thought. In A. Premack, & D. Sperber (Eds.), *Causal cognition*. Oxford: Oxford University Press.

Keil, F. C., & Wilson, R. A. (Eds.). (2000). *Explanation and cognition*. Cambridge, MA: MIT Press.

Kelemen, D. (1999). The scope of teleological thinking in preschool children. *Cognition, 70*, 241–272.

Kelemen, D., Callanan, M., Casler, K., & Pérez-Granados, D. R. (2005). Why things happen: Teleological explanation in parent-child conversations. *Developmental Psychology, 41*, 251–264.

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.
Lombrozo, T. (2006). The structure and function of explanations. *Trends in Cognitive Sciences, 10*, 464–470. http://dx.doi.org/10.1016/j.tics.2006.08.004.
Lombrozo, T. (2009). Explanation and categorization: How "why?" informs "what?". *Cognition, 110*, 248–253.
Lombrozo, T. (2016). Explanatory preferences shape learning and inference. *Trends in Cognitive Sciences, 20*, 748–759. http://dx.doi.org/10.1016/j.tics.2016.08.001.
Macnamara, J. (1986). *A border dispute: The place of logic in psychology.* Cambridge, MA: MIT Press.
Margolis, E. (1998). How to acquire a concept. *Mind and Language, 13*, 347–369.
Markson, L., Diesendruck, G., & Bloom, P. (2008). The shape of thought. *Developmental Science, 11*, 204–208.
Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology, 20*, 158–190.
Moravcsik, J. M. E. (1974). Aristotle on adequate explanations. *Synthese, 28*, 3–17.
Moravcsik, J. M. E. (1975). Aitia as generative factor in Aristotle's philosophy. *Dialogue, 14*, 622–638.
Moravcsik, J. M. E. (1991). What makes reality intelligible? In L. Judson (Ed.). *Aristotle's physics: A collection of essays*. Oxford: Oxford University Press.
Prasada, S. (2016). Mechanisms for thinking about kinds, instances of kinds and kinds of kinds. In D. Barner, & A. S. Baron (Eds.), *Core knowledge and conceptual change*. Oxford University Press.
Prasada, S. (2017). The scope of formal explanation. *Psychonomic Bulletin & Review*.
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.
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.
Salmon, W. C. (1999). *Four decades of scientific explanation.* Minneapolis: University of Minneapolis Press.
Shipley, E. E. (1993). Categories, hierarchies, and induction. *Psychology of Learning and Motivation, 30*, 265–301.
Sloman, S. A., Love, B. C., & Ahn, W. K. (1998). Feature centrality and conceptual coherence. *Cognitive Science, 22*, 189–228.
Strevens, M. (2000). The essentialist aspect of naive theories. *Cognition, 74*, 149–175.
Taylor, M. G., Rhodes, M., & Gelman, S. (2009). Boys will be boys; Cows will be cows: Children's essentialist reasoning about gender categories and animal species. *Child Development, 80*, 461–481.
Wellman, H. M., Hickling, A. K., & Schult, C. A. (1997). Young children's psychological, physical, and biological explanations. In H. M. Wellman, & K. Inagaki (Eds.). *The emergence of core domains of thought: Children's reasoning about physical, psychological and biological phenomena* (pp. 7–25). San Francisco, CA: Jossey-Bass.
Xu, F. (2005). Categories, kinds, and object individuation in infancy. In L. Gershkoff-Stowe, & D. Rakison (Eds.), *Building object categories in developmental time. Papers from the 32nd Carnegie symposium on cognition* (pp. 63–89). New Jersey: Lawrence Erlbaum.
Xu, F. (2012). Count nouns, sortal concepts, and the nature of early words. In F. J. Pelletier (Vol. Ed.), *Kinds, things & stuff: New directions in cognitive science: Vol. 12*. New York, NY: Oxford University Press.