

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

When and How Children Use Explanations to Guide Generalizations

Permalink

<https://escholarship.org/uc/item/2c50g7rj>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors

Vasilyeva, Nadya

Ruggeri, Azzurra

Lombrozo, Tania

Publication Date

2018

When and How Children Use Explanations to Guide Generalizations

Nadya Vasilyeva¹ (vasilyeva@berkeley.edu), Azzurra Ruggeri² (ruggeri@mpib-berlin.mpg.de),
Tania Lombrozo¹ (lombrozo@berkeley.edu)

¹Department of Psychology, University of California, Berkeley, 3210 Tolman Hall, Berkeley, CA 94720 USA

²MPRG iSearch, Max Planck Institute for Human Development, Lentzeallee 94 14195 Berlin, Germany / School of Education,
Technical University Munich, Marsstraße 20-22 80335 Munich, Germany

Abstract

Explanations often highlight inductively rich relationships that support further generalizations: learning that the knife is sharp *because it is for cutting*, we correspondingly infer that other things for cutting might also be sharp. When do children appreciate that explanations are good guides to generalization? We report a study in which 108 4- to 7-year-old children evaluated mechanistic, functional, and categorical explanations for the properties of objects, and subsequently generalized those properties to novel objects on the basis of shared mechanisms, functions, or category membership. Older children, but not younger children, were significantly more likely to generalize when the explanation they had received matched the subsequent basis for generalization (e.g., generalizing on the basis of a shared mechanism after hearing a mechanistic explanation). These findings shed light on how explanation and generalization are coordinated over development, as well as the role of explanations in young children's learning.

Keywords: explanation, generalization, inductive inference, prediction, mechanistic, functional, teleological, categorical, formal

A fidget spinner is an object with a ball bearing that allows it to spin with the flick of a finger, helping the user to relieve nervous energy. For a child encountering this object and its properties for the first time, the task of generalizing from a single example is fraught with peril. For starters, there are multiple relationships to consider: a mechanistic relationship between having a ball bearing and spinning, a functional relationship between spinning and relieving nervous energy, and a categorical (or “formal”) relationship between some of these properties and the category membership of the object. Each of these relationships has the potential to support generalizations to new objects. For instance, other objects might spin in a similar way because they have ball bearings (the mechanistic relationship), because they are for relieving nervous energy (the functional relationship), or because they are fidget spinners (the categorical relationship). How do children navigate this space of possibilities?

In the current paper, we test the hypothesis that explanations help children generalize from known to novel cases by highlighting some relationships as more inductively powerful than others. Specifically, the child who learns that the object spins “because it has a ball bearing” might be more inclined to generalize the property of spinning on the basis of the mechanistic relationship than the child who learns that it spins “because it is a fidget spinner,” no matter that both children know that all three relationships (mechanistic, functional, and categorical) in fact hold for the object being explained.

This example suggests that when we explain, we do more than identify true claims that happen to be valid in a particular case: we highlight generalizable patterns that can

be extended beyond the particular observation being explained. Indeed, this idea has motivated accounts of explanation in both philosophy and psychology (Heider, 1958; Quine & Ullian, 1970). According to the “Explanation for Export” proposal developed in Lombrozo and Carey (2006), explanations serve the cognitive function of supporting generalization beyond the specific case being explained. Indeed, empirical work with adults supports the idea that explanations guide generalization. Sloman (1994) found that whether adults generalize properties from known to novel cases depends on how they explain the presence of those properties. For instance, participants gave a high rating to the claim that “secretaries have a hard time financing a house” given that “furniture movers have a hard time financing a house” because both claims support the same explanation: low wages. Ratings were lower when the claims involved different explanations – for instance, “secretaries have bad backs” (from sedentary work) given that “furniture movers have bad backs” (from heavy lifting).

Along this line, Vasilyeva and Coley (2013) found that an individual's explanation for a property predicted subsequent generalizations of that property, with the additional finding that different *types* of explanations corresponded to different *types* of generalizations. In their study, participants were prompted to generalize the properties of living things (e.g., if substance B6 is found in ducks, what else is likely to have substance B6?). In doing so, participants often generated spontaneous explanations for why the property held. When participants generated category-based explanations (e.g., ducks have it because “it's a bird thing”) or functional explanations (e.g., because “it protects them from the cold”), they tended to project the property to targets related categorically rather than ecologically (e.g., to “other birds” rather than “their predators”). However, this tendency was flipped for mechanistic explanations (e.g., “they got it from their food”). In an experimental task, Lombrozo and Gwynne (2014) found that participants who received an explanation that was either mechanistic or functional tended to favor subsequent generalizations that preserved the kind of relationship featured in the explanation.

A recent study by Vasilyeva, Wilkenfeld, and Lombrozo (2017) illustrates the reverse direction of influence: the types of generalizations that participants anticipated making affected their evaluations of explanations. Participants were led to expect that they would make later inferences about the presence or absence of a target property based on information about causal mechanisms, functions, or category membership. This manipulation in turn affected how highly participants rated mechanistic, functional, and categorical explanations, respectively, suggesting that participants

avored explanations based on the kind of generalization that would support their subsequent judgments.

Although there is good evidence that adults can use explanations of different types to guide generalization, it is unclear when this ability emerges and how it develops. Do young children coordinate explanations and generalizations from the moment they appreciate the distinctions between different kinds of explanations? Or is this an ability that develops more gradually, as children learn that explanations are related to inductive potential? In the current paper, we investigate whether 4- to 7-year-old children are more likely to generalize a property on the basis of a given feature (a causal mechanism, function, or category membership) after they previously heard an explanation underwritten by a generalization congruent with that feature (i.e., mechanistic, functional, or categorical).

Examining how the coordination between explanation and generalization develops is important because it offers insights into one of the most basic questions in cognitive development: how children learn so much from limited input. Specifically, how do children constrain the range of hypotheses they entertain when generalizing from the known to the unknown? Receiving explanations from adults or peers could be one source of constraint, directing children to the generalizations that are most likely to be useful. Indeed, a large literature suggests that seeking, generating, and receiving explanations is a powerful basis for learning (e.g., Chi, de Leeuw, Chiu, & LaVancher, 1994; see Lombrozo, 2012 and Wellman, 2011, for reviews). The link between explanations and inductive potential could be one reason explanation has such powerful effects.

The development of explanation and generalization

Children's early explanatory sophistication suggests that the coordination between explanation and generalization could be in place by age 4 or 5. Children as young as 2 years of age provide and seek explanations (Callanan & Oakes, 1992; Hickling & Wellman, 2001), and they begin to differentiate explanations from non-explanations by age 3 (Frazier, Gelman, & Wellman, 2009). There is also evidence that explanation is coordinated with other representations and judgments from an early age. Three- to 4-year-olds invoke domain-appropriate mechanisms in their explanations, for example explaining biological phenomena by appeal to biological factors (Hickling & Wellman, 2001; Inagaki & Hatano, 2002; Schult & Wellman, 1997).

Prior work also suggests that young children differentiate between different *types* of explanations (mechanistic, functional, and categorical). By age 4, children request explanations of different types in different domains (Greif, Kemler Nelson, Keil, & Gutierrez, 2006), successfully generalize the form of an explanation (i.e., mechanistic versus functional) to novel cases (Lombrozo, Bonawitz, & Scalise, 2018), and appreciate when formal (categorical) explanations are appropriate (Haward, Wagner, Carey, & Prasada, In Prep). There is also evidence that through grade school, children favor functional over mechanistic explanations (Kelemen, 1999; Kelemen & DiYanni, 2005;

but see Keil, 1994, 1996; Lombrozo, Bonawitz, & Scalise, 2018).

Finally, there is a great deal of evidence demonstrating that receiving and generating explanations can affect learning and subsequent inferences (e.g., Wellman & Lagattuta, 2004; Walker, Lombrozo, Legare, & Gopnik, 2014; Walker, Lombrozo, Williams, Rafferty, & Gopnik, 2017). Repeatedly prompting 3-year-old children to explain a character's behavior in a false-belief vignette, for example, can accelerate their understanding of theory of mind (e.g., Amsterlaw & Wellman, 2006). As another example, Walker et al. (2014) found that prompting 3- to 5-year-old children to explain why objects made a machine go led them to privilege invisible properties and category membership, as opposed to appearance, as a basis for subsequently generalizing a causal property.

Despite this evidence of explanatory sophistication in the preschool years, a handful of findings suggest that explanation is somewhat quarantined from prediction, at least until age 4. Specifically, several studies have found an "explanation advantage," such that successful explanation precedes accurate prediction. For instance, children are able to explain why someone did not choose to eat a contaminated food before they can predict the same event (Amsterlaw & Wellman, 2006; Bartsch & Wellman, 1989; Legare, Wellman, & Gelman, 2009). Wellman (2011) explains this progression in terms of the difference between *postdiction* and *prediction*: in the former case an additional piece of information (the actual outcome) is known. As a form of postdiction, explanation could involve a lower cognitive burden and serve as a stepping stone to later prediction.

By age 8, children are capable of generating sophisticated and context-appropriate explanations and predictions, but they do not yet show adult-like coordination between explanations and generalization. For example, one child in Vasilyeva and Coley (In Prep) explained why a zebra and savannah grass are both "sick" with the same disease by appeal to causal transmission from grass to zebra: "because it eats grass." Yet when asked what else might be sick with the same disease, the child relied on a completely independent basis for generalization, saying "a werewolf, because both zebras and werewolves have a tail." (Admittedly, this was a demanding paradigm that required children to generate their own explanations and predictions from prior domain knowledge, so it could be that task demands masked existing competence.)

Summing up this research, there are good reasons to expect that even young children can effectively learn from explanations, suggesting some level of coordination between explanation and subsequent generalizations. At the same time, the coordination of explanation and prediction is tenuous at age 4 (as reflected in the "explanation advantage" and other evidence of explanation-prediction asymmetries, see Nancekivell & Friedman, 2017), and still fragile at age 8 (Vasilyeva & Coley, in prep). In the present study, we thus focused on 4- to 7-year-olds as an age range during which we might expect to see developmental change in the coordination of explanation and generalization.

Experiment

The goal of this study was to investigate when and how children begin to use explanations as a guide to subsequent generalizations. Children evaluated a mechanistic, functional, or categorical explanation for the property of a novel object, and then guessed whether that property would generalize to a hidden object, knowing only that it shared a mechanism, function, or category membership with the initial object.

The main manipulation concerned the “congruence” between the explanation that was presented, on the one hand, and the feature that was shared by the two objects, on the other. On *congruent* trials, the same feature that was invoked in the explanation of the property was shared by both objects. For example, a participant might hear a mechanistic explanation (e.g., “Why is it sticky? It is sticky because it has a special tape on one side”), and then be presented with a hidden object sharing the same mechanism invoked in the explanation (“it has a special tape on one side...do you think it is sticky?”). On *incongruent* trials, the explanation mentioned one feature (e.g., mechanistic, “because it has a special tape”), but the hidden object shared a different feature (e.g., functional, “it can pick up marbles”).

We tested children aged 4 to 7, with the expectation that with age children would become increasingly likely to use explanations to guide subsequent generalizations, as reflected in higher generalization ratings on congruent versus incongruent trials. The study design also allowed us to examine whether effective coordination depends on the *type* of explanatory relationship involved (i.e., mechanistic, functional, or categorical). Coordination might emerge earlier for the kinds of relationships that are most widely applicable (arguably mechanistic), for those that figure in favored kinds of explanations (arguably functional; Kelemen, 1999, 2005), or for those that dominate children’s early generalizations (arguably categorical; Coley, 2012).

Method

Participants We recruited 54 4- and 5-year-olds (mean age 5 years 2 months; range 4.00-5.98 yrs) and 54 6- and 7-year-olds (mean age 6 years 11 months; range 6.00-7.96 yrs) from local museums and preschools. Data from an additional eight children were not included due to a failure of video-recording equipment (six children) or experimental error (two children).

Materials, Design and Procedure

Each participant completed three trials. Each trial involved a learning phase, an explanation phase, and a generalization phase (see Table 1).

In the *learning phase*, the experimenter introduced a novel type of object (base object; e.g., a “dax”) and presented three features in the form of a causal chain (e.g., has special tape on one side → is sticky → can pick up marbles). Each object and each feature was represented with a color illustration printed on a laminated card; the feature cards were laid out one by one, illustrating each feature as it was introduced. The causal chain was then repeated once.


In the *explanation phase*, the experimenter laid out a black-and-white silhouette of the same type of object, and asked a

why-question about the middle feature in the causal chain (e.g., “Why is it sticky?”), addressing a puppet on a laptop screen. In a short video clip, the puppet provided an explanation. The child was then asked to evaluate the explanation using a two-step, four-point thumb scale ranging from “really bad” to “really good.” After the explanation-evaluation, the experimenter repeated the explanation and removed all the pictures from the table.

In the *generalization phase*, the experimenter presented a closed box tied with a ribbon. On top of the box was a transparent pocket with a face-down card representing an object feature. The experimenter then said that Julia (another puppet) wanted to know if the object in the box had a certain feature (always the middle feature from the causal chain, e.g., being “sticky”). The experimenter stated that she did not know the answer, but that they could check what the box said. The experimenter flipped the card on top of the box to reveal a picture representing one of the features (either a mechanism, function, or category membership feature), and asked the child to make a guess (“Do you think it is sticky or not sticky?”). This was followed by a further rating (“For sure [sticky / not sticky], or maybe [sticky / not sticky]”) to obtain a 4-point rating. Then the box was removed, and the experimenter moved on to the next trial.

Table 1. Sample script from an incongruent trial (mechanistic explanation, function-based generalization). Arrows were not presented.


Learning phase:



This is a dax. A dax is a kind of toy.
Let me tell you a few things about daxes.

They have special tape on one side... → ...which makes them sticky. The sticky side.. → ...makes it easy to pick up marbles.

Explanation phase:



Here's one. Why is it sticky? Let's ask my friend Mike. Mike sometimes says things that are smart, and sometimes says things that are silly. Let's see what Mike thinks! Mike, why is this sticky?

Because it has special tape on one side!

Mike says it is sticky because it has special tape on one side.

[Explanation rating]
What do you think about Mike's explanation?
Is it a good explanation or is it a bad explanation?
Is it really [good/bad] or is it kinda [good/bad]?
Ok, remember, Mike said it is sticky because it has special tape on one side.

Generalization phase:



Now, There is something in this box. Julia wants to know if it is sticky. I don't know. But let's see what the box says.
[Experimenter flips the label on top]. It says that it *can pick up marbles*.
[Generalization rating] Do you think it is sticky or not sticky? □ Is it for sure (not) sticky or maybe (not) sticky?

Table 2. Mechanistic, functional, and categorical explanations.

Why-Question	Mechanistic explanation	Functional explanation	Categorical explanation
Why is this sticky?	Because it has special tape on one side.	Because that way it's easy to pick up marbles.	Because it is a dax.
Why does this glow in the dark?	Because it is covered with special paint.	Because that way it's easy to spot it at night.	Because it is a zark.
Why is this bouncy?	Because it has a spring on one side.	Because that way it can go higher.	Because it is a wug.

Trials varied in the object introduced in the learning phase, in the explanations offered in the explanation phase, and in the generalizations solicited in the generalization phase. The explanations were either mechanistic (citing the preceding feature in the causal chain), functional (citing the final feature in the causal chain), or categorical (citing category membership). Across trials, each participant evaluated one explanation of each type.

The solicited generalizations were based on a shared mechanism feature, function feature, or category membership feature, and they were coded as either “congruent” (on one trial) or “incongruent” (on two trials). On congruent trials, the generalization target and the base object shared the same feature that was invoked in the preceding explanation. For instance, if a functional explanation was offered in the explanation phase, a congruent generalization trial would involve generalizing from the presence of the function feature. On the remaining two trials, participants were presented with “incongruent generalization” tasks, combining an explanation of one type and a generalization task of a different type (e.g., mechanism-based generalization following a categorical explanation).

The pairing between the three objects, three explanation types, and three generalization types was counterbalanced using a Latin square design, producing nine unique conditions. Across participants, each explanation type was thus paired with each object and with each generalization type.

Results

Children’s explanation ratings were recoded onto a four-point scale (corresponding to “really bad,” “kind of bad,” “kind of good,” and “really good”). The generalization ratings were also recoded onto a four-point scale (corresponding to “for sure does not have,” “maybe does not have,” “maybe has,” and “for sure has” the target property).

Explanation evaluation. A 2 (age group: younger, older) $\times$ 3 (explanation type: mechanistic, functional, categorical) mixed ANOVA on explanation evaluation ratings¹ revealed a main effect of age group, $F(1,105)=13.00$, $p<.001$, $\eta^2=.110$: younger children gave higher ratings ($M=3.45$) than older children ($M=3.06$). There was also a main effect of explanation type, $F(2,210)=22.08$, $p<.001$, $\eta^2=.174$: mechanistic explanations ($M=3.56$) were rated higher than functional explanations ($M=3.33$, $p=.027$), which were rated higher than categorical explanations ($M=2.88$, $p<.001$; causal

¹ Due to an experimental error or audio/video equipment failure, data from one explanation trial and four generalization trials were lost (from two younger and two older children). For a given analysis, participants with any missing data were excluded.




Figure 1: Mean explanation evaluation ratings as a function of explanation type and age group. Error bars correspond to 1 SEM; $*p < .05$.

vs. categorical $p<.001$). However, these effects were qualified by a significant interaction (see Figure 1), $F(2,210)=4.23$, $p=.016$, $\eta^2=.039$. Older children differentiated among all three explanation types, rating mechanistic explanations higher than functional ($p=.033$), and functional higher than categorical ($p<.001$). Younger children rated mechanistic explanations higher than categorical ($p=.009$), but their ratings of functional explanations did not differ from either mechanistic ($p=.112$) or categorical explanations ($p=.316$).

Additional one-sample t-tests showed that both age groups rated all explanations above the scale midpoint (all p 's $>.001$), except the older children’s ratings of categorical explanations ($M=2.5$, $p>.999$).

Property generalization. Collapsing across the initial explanations, a 2 (age group: younger, older) $\times$ 3 (generalization type: cause-based, function-based, category-based) mixed ANOVA on generalization ratings revealed no significant effects. This suggests that the stimuli were well-matched across generalization types, providing an even playing field on top of which the preceding explanations might exert some effect. Additional one-sample t-tests showed that both age groups rated all generalizations above the scale midpoint (all p 's $>.001$).

Relationship between explanation and property generalization. To test for a relationship between rated explanations and subsequent generalizations, we first examined whether generalization ratings were reliably higher for congruent trials relative to incongruent trials.

Generalization ratings were first analyzed in a 2 (congruence: congruent trial, incongruent trial) $\times$ 2 (age group: younger, older) mixed ANOVA. The analysis revealed a main effect of congruence, $F(1,106)=9.87$, $p=.002$, $\eta^2=.085$, with higher generalization ratings on congruent trials ($M=3.63$) than



Figure 2: Mean generalization ratings as a function of congruence and age group. Error bars correspond to 1 SEM; $*p < .05$.




Figure 3. Mean generalization ratings as a function of age group, congruence, and generalization type. Error bars correspond to 1 SEM.

on incongruent trials ($M=3.44$). This effect was qualified by a marginal interaction with age group, $F(1,106)=3.78$, $p=.054$, $\eta_p^2=.034$. Planned contrasts revealed a developmental change in the relationship between explanation and generalization: the older children rated congruent generalizations significantly higher than incongruent generalizations, $p<.001$; the younger children did not show a significant effect of congruence, $p=.399$ (see Figure 2).

To evaluate whether some kinds of explanations fostered congruent generalizations more effectively than others (or equivalently, whether some kinds of generalizations were more susceptible to the effect of a congruent explanation), we conducted three 2 (congruence: congruent, incongruent) $\times$ 2 (age group: younger, older) ANOVAs, one for each kind of generalization (mechanism-, function-, or category-based; see Figure 3).

Mechanism-based generalization was affected by congruence, $F(1,102)=4.57$, $p=.035$, $\eta_p^2=.043$; this effect was qualified by a marginal interaction with age, $F(1,102)=3.77$, $p=.055$, $\eta_p^2=.036$. The older children ($p=.003$), but not the younger children ($p=.809$), favored congruent generalizations.

Function-based generalization was likewise influenced by congruence, $F(1,101)=5.32$, $p=.023$, $\eta_p^2=.050$. Although the interaction with age was not significant, $F(1,101)=.64$, $p=.426$, we conducted planned pairwise comparisons; again, only the older children showed a significant effect of congruence ($p_{old}=.029$, $p_{young}=.293$).

Category-based generalization was not affected by congruence or by age group, all $ps \geq .567$. Finally, the main effect of age group was not significant in any of the analyses reported in this section, $ps \geq .125$.

Effects of explanation quality. The preceding analyses suggest that for older children, congruent explanations indeed guide subsequent generalization. We next investigated whether this effect was moderated by the perceived quality of the rated explanation. That is, were children more likely to generalize a property on the basis of some feature when they found an explanation that appealed to that feature good versus bad? For this analysis, we recoded explanation ratings

as “high quality” if an explanation received the highest possible rating of 4, and as “low quality” otherwise. Older children were significantly more likely to generalize on congruent trials if they rated the preceding explanation as “high quality” than “low quality” ($M_{high}=3.83$, $M_{low}=3.48$, $t(49.26)=2.34$, $p=.023$, Cohen’s $d=.62$). In contrast, the younger group showed no relationship between their ratings of explanation quality and subsequent generalizations ($M_{high}=3.61$, $M_{low}=3.75$, $t(52)=-.66$, $p=.515$).

Discussion

The primary aim of this study was to document the emergence of coordination between explanation and generalization. We found that by age 6-7, children were significantly more likely to generalize a property on the basis of some feature if they had previously heard an explanation for that property appealing to the same feature. Moreover, the effect of the explanation was greater for children who judged it to be a *good* explanation. Neither of these effects was found for children aged 4-5.

We also found a developmental shift in how strongly children differentiated explanations of different kinds. While 6- and 7-year olds reliably favored mechanistic explanations over functional explanations, and functional explanations over categorical explanations, the preferences of 4- and 5-year-olds were less robust.

We also found potential differences in the relationship between explanation and generalization depending on the type of relationship involved. By age 6-7, mechanistic and functional explanations reliably boosted congruent generalizations, but there was no such relationship between categorical explanation and category-based generalization. The latter is surprising given evidence that young children readily engage in category-based inference when category labels are provided (Gelman & Coley, 1990), and that adults can successfully coordinate categorical explanations and generalizations (Vasilyeva & Coley, 2013; Vasilyeva, Wilkenfeld, & Lombrozo, 2017). Along with the finding that older children’s coordination between explanation and generalization was moderated by perceived explanation quality, the moderating effect of explanation type speaks against the idea that children were simply following a low-level strategy of providing higher ratings when a generalization “matched” what came before.

Our findings prompt a variety of follow-up questions. What changes between ages 4-5 and 6-7, such that we see greater coordination between explanation and generalization alongside greater differentiation between different types of explanations? Are we right to suggest that congruent explanations promote generalization, or could it be that incongruent explanations *suppress* generalization?

Another open question concerns the role of pedagogy in our task. Rather than presenting children with an explanation from an authoritative source, children in our study were presented with an explanation from a puppet who “sometimes says things that are silly,” and they were then asked to evaluate the explanation themselves. It is plausible that presenting an expert explanation in a pedagogical context would have a stronger effect on subsequent

² We originally intended to perform a median split of explanation ratings, but the median for both age groups was 4, so we used the next closest split between ratings of 3 and 4. Splitting the scale in the middle (between 2 and 3) produced the same pattern of results, but also resulted in highly uneven group sizes (8 vs. 46, 14 vs. 40), making the statistical tests less reliable.

generalizations. It's also possible that such explanations would not only encourage some generalizations, but actively restrict others, much like pedagogical demonstrations can discourage children from more open-ended exploration (Bonawitz, Shafto, Gweon, Spelke, & Goodman, 2011). This could be the “dark side” of explanation's positive inductive role: by encouraging learners to favor some inductive hypotheses over others, an explanation could restrict some forms of exploration or inference (see also Legare, 2012).

In sum, our findings capture an important developmental transition: from generalization unconstrained by explanation, to generalization *guided* by it. As children master the coordination of explanation and generalization, they can benefit from the advantages enjoyed by adults, including the ability to effectively constrain a range of inductive hypotheses to the most relevant and plausible subset, and to generate flexible inferences from the same observation depending on one's context (Vasilyeva & Coley, 2013) and goals (Vasilyeva, Wilkenfeld, & Lombrozo, 2017). At the same time, greater constraint does not come without costs – children could lose some of the advantages of unconstrained exploration and an undifferentiated hypothesis space. To borrow a metaphor from computer science, explanation could guide the transition from “low-temperature search” to “high-temperature search” (Gopnik, Griffiths, & Lucas, 2015), driven by a growing appreciation for the high inductive potential of the relationships that good explanations extract from the world.

Acknowledgments

Varieties of Understanding project funded by the John Templeton Foundation; McDonnell Foundation Award in Understanding Human Cognition; both awarded to TL.

References

- Amsterlaw, J., & Wellman, H. M. (2006). Theories of mind in transition: A microgenetic study of the development of false belief understanding. *Journal of cognition and development*, 7(2), 139-172.
- Bartsch, K., & Wellman, H. (1989). Young children's attribution of action to beliefs and desires. *Child development*, 946-964.
- Bonawitz, E., Shafto, P., Gweon, H., Goodman, N., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, 120(3), 322-330.
- Callanan, M. A., & Oakes, L. M. (1992). Preschoolers' questions and parents' explanations: Causal thinking in everyday activity. *Cognitive Development*, 7(2), 213-233.
- Chi, M. T., De Leeuw, N., Chiu, M. H., & LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive science*, 18(3), 439-477.
- Coley, J.D. (2012). Where the wild things are: Informal experience and ecological reasoning. *Child Development*, 8(2), 992-1006.
- 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, S.A., & Coley, J.D. (1990). The importance of knowing dodo is a bird: Categories and inferences in 2-year-old children. *Developmental Psychology*, 26(5), 796-804.
- Gopnik, A., Griffiths, T. L., & Lucas, C. G. (2015). When younger learners can be better (or at least more open-minded) than older ones. *Current Directions in Psychological Science*, 24(2), 87-92.
- Greif, M. L., Kemler Nelson, D. G., Keil, F. C., & Gutierrez, F. (2006). What do children want to know about animals and artifacts? Domain-specific requests for information. *Psychological Science*, 17(6), 455-459.
- Haward, P., Wagner, L., Carey, S., & Prasada, S. (In Prep). The development of principled connections and kind representations.
- Heider, F. (1958). *The psychology of interpersonal relations*. Hillsdale, NJ: Psychology Press.
- Hickling, A. K., & Wellman, H. M. (2001). The emergence of children's causal explanations and theories: Evidence from everyday conversation. *Developmental Psychology*, 37(5), 668.
- Inagaki, K., & Hatano, G. (2006). Young children's conception of the biological world. *Current Dir. in Psych. Science*, 15(4), 177-181.
- Keil, F.C. (1994). The birth and nurturance of concepts by domains: The origins of concepts of living things. In *Mapping the Mind: Domain Specificity in Cognition and Culture* (Eds. L.A. Hirschfeld and S.A. Gelman). Cambridge: CUP.
- Keil, F.C. (1996). The growth of causal understandings of natural kinds. In *Causal Cognition: A Multidisciplinary Debate* (Eds. D. Sperber, D. Premack, & A.J. Premack). New York: OUP.
- Kelemen, D. (1999). The scope of teleological thinking in preschool children. *Cognition*, 70(3), 241-272.
- Kelemen, D., & DiYanni, C. (2005). Intuitions about origins: Purpose and intelligent design in children's reasoning about nature. *Journal of Cognition and Development*, 6(1), 3-31.
- Legare, C. H. (2012). Exploring explanation: Explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child development*, 83(1), 173-185.
- Legare, C. H., Wellman, H. M., & Gelman, S. A. (2009). Evidence for an explanation advantage in naïve biological reasoning. *Cognitive psychology*, 58(2), 177-194.
- Lombrozo, T. (2012). Explanation and abductive inference. *Oxford handbook of thinking and reasoning*, 260-276.
- Lombrozo, T., Bonawitz, E. B., & Scalise, N. R. (2018). Young children's learning and generalization of teleological and mechanistic explanations. *Journal of Cognition and Development*.
- Lombrozo, T., & Carey, S. (2006). Functional explanation and the function of explanation. *Cognition*, 99(2), 167-204.
- Lombrozo, T., & Gwynne, N. Z. (2014). Explanation and inference: mechanistic and functional explanations guide property generalization. *Frontiers in human neuroscience*, 8, 700.
- Nancekivell, S.E., & Friedman, O. (2017). She bought the unicorn from the pet store: Six-to-seven-year-olds are strongly inclined to generate natural explanations. *Developmental psychology*, 53(6), 1079.
- Quine, W. V., & Ullian, J. S. (1970). The web of belief. *New York*.
- Prasada, S., & Dillingham, E. M. (2009). Representation of principled connections: A window onto the formal aspect of common sense conception. *Cognitive Science*, 33(3), 401-448.
- Schult, C. A., & Wellman, H. M. (1997). Explaining human movements and actions: Children's understanding of the limits of psychological explanation. *Cognition*, 62(3), 291-324.
- Sloman, S. (1994). When explanations compete: The role of explanatory coherence on judgements of likelihood. *Cognition*, 52, 1-21.
- Vasilyeva, N., & Coley, J. (2013, January). Evaluating two mechanisms of flexible induction: selective memory retrieval and evidence explanation. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 35, No. 35).
- Vasilyeva, N. & Coley, J.D. (In Prep). Explanation and generalization.
- Vasilyeva, N., Wilkenfeld, D., & Lombrozo, T. (2017). Contextual utility affects the perceived quality of explanations. *Psychonomic bulletin & review*, 24(5), 1436-1450.
- Walker, C. M., Lombrozo, T., Legare, C. H., & Gopnik, A. (2014). Explaining prompts children to privilege inductively rich properties. *Cognition*, 133(2), 343-357.
- Walker, C. M., Lombrozo, T., Williams, J. J., Rafferty, A. N., & Gopnik, A. (2017). Explaining constrains causal learning in childhood. *Child development*, 88(1), 229-246.
- Wellman, H. M. (2011). Reinvigorating explanations for the study of early cognitive development. *Child Development Perspectives*, 5(1), 33-38.
- Wellman, H. M., & Lagattuta, K. H. (2004). Theory of mind for learning and teaching: The nature and role of explanation. *Cognitive development*, 19(4), 479-497.