

Mechanistic Knowledge Generalizes Differentially

Aaron Chuey, Mark Sheskin, Frank Keil
(aaron.chuey@yale.edu, mark.sheskin@yale.edu,
frank.keil@yale.edu)

Yale Department of Psychology, 2 Hillhouse Ave. New Haven, CT 06511 USA

Abstract: When inferring the extent of others' knowledge from samples of what they know, certain kinds of samples imply richer content. One candidate kind is knowledge of causal mechanism. In the current study, we investigate whether children and adults think that knowledge about mechanism generalizes more broadly than non-mechanistic factual knowledge. We find an early-emerging assumption that mechanistic knowledge about a basic level category implies greater knowledge about a superordinate category, compared to factual knowledge about the same basic level category. Even young children have a sophisticated sense of how causal mechanisms generalize across categories, despite possessing little mechanistic knowledge themselves. These intuitions likely support the epistemic inferences we make from early childhood onward.

Keywords: mechanism, causal reasoning, knowledge, category learning, epistemic inference

Introduction

To benefit from a world full of information, we must make a variety of inferences about who possesses useful information and how much of it they possess. To do so, we also need a sense of how broadly knowledge and information generalize. For example, if someone knows how tractor engines work, should I assume that she is also knowledgeable about cars? Airplanes? Iguanas? Even children demonstrate sophisticated intuitions about how knowledge (Keil et al, 2008) and explanations (Johnston et al, 2017) generalize across kinds and domains, but it is unclear what underlies these intuitions. Here, we argue that a sense of shared understandings of causal mechanisms plays an important role in the way we generalize knowledge from one kind to another, and that this sense is evident in children's early emerging epistemic intuitions.

A growing literature supports the importance of mechanism in young children's strategies for evaluating and structuring knowledge. For example, preschoolers judge that someone who can fix an object has more causal knowledge about it than someone who knows its name (Kushnir Vredenburgh, & Schneider, 2013). Because children associate mechanistic knowledge with an ability to fix (Lockhart et al, under review), intuitions about causal mechanisms may influence children's expectations about who possesses useful information.

Young children also show an appreciation for mechanism when reasoning about the way knowledge is structured beyond particular knowers. By age 5, children group biological and psychological processes separately based on

a notion of shared causal mechanisms (Erickson, Keil & Lockhart, 2010). As they get older, children develop a stratified sense of difficulty for the sciences (Keil, Lockhart, & Schlegel, 2010), suggesting that intuitions about causal mechanisms affect not only the way children organize knowledge, but also their attitudes towards it.

Causal mechanism also pervades children's explanatory preferences. When requesting information, young children are often not satisfied with statements of fact or circular reasons, instead preferring causal explanations (Corriveau & Kurkul, 2014; Frazier, Gelman & Wellman, 2009; 2016). Children's desire for rich information increases with age; requests for causally rich explanations take up an increasingly large proportion of children's questions as they reach elementary school (Chouinard et al., 2007). For example, one study found that "how" questions make up only 3.5% of 3-year-olds' questions, but 19.8% of 5-year olds' questions (Callanan & Oakes, 1992). Furthermore, young children remember a larger number of causally relevant features when they explain phenomena, as opposed to merely reporting on them (Legare & Lombrozo, 2014; Walker et al, 2017). Preschoolers likewise privilege "deep" properties in their own explanations, favoring more inductively powerful features of a system when engaging in explanation (Walker et al, 2014).

Although there is clear evidence that notions of mechanism influence children's reasoning about knowledge and explanations, what accounts for its utility? Here, we argue that one critical feature is mechanism's generalizability across related kinds. Trouche et al. (2017) found that mechanism-focused instruction shifted elementary school children's complexity intuitions more than factual non-mechanistic instruction, suggesting that mechanism cues children to certain properties that other forms of factual information do not. Children's mechanism-induced complexity intuitions also propagate to related entities (Trouche et al, under review), suggesting that children may expect knowledge of mechanism-related properties to apply broadly. For example, if someone knows how a car works, we might expect that knowledge to apply, at least in part, to tractors because most vehicles work in a broadly similar way utilizing broadly similar internal components. Here we investigate whether children think that mechanistic knowledge about basic level categories (e.g. cars, clocks, and smartphones) generalizes to knowledge about their superordinate level categories (e.g. vehicles, machines, and electronics).

In the current study, children (ages 6 to 9 years old) and adults were presented with two twins, one possessing mechanistic knowledge about a basic level category (e.g., clocks) and another possessing factual non-mechanistic knowledge about that category. Participants were then asked which twin knew more about its superordinate category (e.g., machines), a subordinate category (e.g., grandfather clocks), and an unrelated basic level category (e.g., tulips). A superordinate category was included as the key measure: differentially generalizing mechanistic knowledge from a

basic level category to a superordinate level category would suggest that mechanistic knowledge applies more broadly to kinds within that superordinate category compared to non-mechanistic knowledge about the same basic level category. A subordinate category was included to assess the scope of the generalization inferences: differentially generalizing mechanistic knowledge from a basic level category to a subordinate level category would imply that mechanistic knowledge applies to specific kinds within that basic level category, an inference that might appeal to young children if they see mechanistic knowledge as generalizing more strongly to all instances of the category. Not selectively generalizing mechanistic knowledge to the subordinate level would reveal a more nuanced view about mechanistic knowledge's generalizability. An unrelated basic level category was included as a control: not generalizing mechanistic knowledge to an unrelated category ensures that participants did not perceive the mechanistic knower as more intelligent or knowledgeable in general.

Our hypotheses were:

1. Participants across all ages would judge the twin possessing mechanistic knowledge as more knowledgeable about the superordinate category than the twin possessing factual non-mechanistic knowledge, showing that mechanistic knowledge about a basic level category is generalized to its superordinate level category more than factual knowledge.

2. Participants across all ages would not judge the mechanistic twin to know significantly more about the unrelated basic level category. This result would clarify evidence in support of our first hypothesis, demonstrating that mechanistic knowledge does not simply imply greater knowledge in general, but rather about the superordinate category in particular, and also that participants do not perceive the mechanistic twin as merely more intelligent or knowledgeable in general.

3. Participants across all ages would not judge the mechanistic twin to know significantly more about the subordinate level category. This result would further clarify evidence in favor of our first hypothesis, demonstrating that participants perceived the mechanistic and non-mechanistic twins as possessing approximately the same amount of knowledge about categories below the basic level category.

Experiment

Method

Participants Twenty-one 6- and 7-year-olds (11 male, mean age 85 months) and twenty 8- and 9-year-olds (8 male, mean age 106 months) participated in the experiment via TheChildLab.com online platform (Sheskin & Keil, under review). On this platform, researchers can engage in online videoconferences with participants on a web-enabled device. The study stimuli are presented as a PowerPoint presentation shared within the videoconference, and sessions begin with simple warm-up activities (e.g. following a ball through a tube), established as easy for

most participants in previous research. Thirty-nine adults participated in the experiment via Amazon Mechanical Turk for \$.50 payment; thirty-one adults (20 male, mean age 32 years) passed all attention checks and were included in the final sample.

Materials Three stimulus categories were used, with each matched to a different superordinate category, subordinate category, and unrelated basic level category. The three stimulus/superordinate/subordinate/unrelated sets were clocks, machines, grandfather clocks, tulips; cars, (wheeled) vehicles, racecars, sharks; smartphones, electronics, iPhones, tigers. Each category was presented with an image depicting the category, consisting of 6 category exemplars in a white square (to emphasize kind rather than token). These categories were chosen because they represented a broad sample of artifacts familiar to most children.

Design Each question was focused on which one of two cartoon children knew more about a kind. Each pair were introduced as twins, and looked nearly identical except one twin wore blue clothes and the other wore green clothes. The twins were referred to by the color of their clothes (i.e., as "Blue" or as "Green"). The blue twin was always described first for each stimulus category, but the blue twin's knowledge type (mechanistic or non-mechanistic) was counterbalanced across participants. The test categories (superordinate, subordinate, side) were presented in a consistent order for a given participant across all three stimulus items, but order was counterbalanced with either the superordinate category being presented first and the subordinate category last, or vice-versa. The order of the stimulus items was randomized across participants. The study took approximately 8 minutes for children and 5 minutes for adults.

Procedure At the start of the activity, children were presented with a training example that introduced the concept of a yellow equal sign, which would be used in the activity. They were then trained on how to give an answer in the activity, saying "blue" if they chose the blue twin (always on the left), "green" if they chose the green twin (always on the right), and "yellow" if something applied to both twins the same (the yellow equals sign was always shown between the twins). A "same" choice was included because it could reflect a genuine preference, especially for the unrelated category. Adults did not complete this training and were instead instructed "in this survey, you are going to hear about pairs of twins who both read a book about the same topic, so they both learn a lot of information about the same thing. However, the books they read are different, so they each learn different information about the same thing. You will hear about the kind of things that each twin learns, and then your job is to decide who knows more about some different things." Participants were then introduced to both twins and told that each twin read a corresponding colored book, which were identical besides the color of their covers. Both books were about the same stimulus category, so they both learned lots of things about that category, but the books were different so each twin learned different kinds of things

about it. Participants were told one twin learned about how the category works (the mechanistic twin), and the other twin learned facts about the category (the non-mechanistic twin). They were also given two examples of each twin’s mechanistic or non-mechanistic knowledge (see Table 1). Participants were then presented with a test category and asked which twin knew more about the category, or whether they thought the twins knew the same amount about it. For example, children might be asked: “Here are some machines. Who do you think knows more about machines? Blue who knows about how clocks work, Green who knows facts about clocks, or Yellow, do you think they know about the same amount?” Children were asked about each test category in sequence, with each test category consisting a different slide. Adults were presented with all three test categories in sequence on the same page. Participants completed this procedure for each stimulus category.

Stimulus	Mechanistic Knowledge	Non-mechanistic Knowledge
Clocks	For example, she learned what makes the parts of the clock move. As another example, she learned how clocks can keep working for years without stopping.	For example, she learned where clocks were first invented. As another example, she learned how many clocks are made every year.
Cars	For example, he learned how car engines make the car move. As another example, he learned how cars’ brakes make the tires stop spinning.	For example, he learned where the first car engines were built. As another example, he learned how many different companies make tires for cars.
Smartphones	For example, he learned how smartphones’ screens recognize your fingerprint. As another example, he learned how smartphones are able to make many different kinds of sounds.	For example, he learned what kinds of glass smartphones’ screens are made out of. As another example, he learned how many different ringtones are available for smartphones.

Table 1: Knowledge examples

Results

No children were excluded, yielding a sample of 21 6-7 year olds (age group 1), 20 8-9 year olds (age group 2), and 31 adults (age group 3). For purposes of analysis, choosing the mechanistic twin was coded as 1, the non-mechanistic twin as -1, and knowing the same as 0. Scores were aggregated across stimulus items for each categorical level, yielding a superordinate level score, a subordinate level score, and an unrelated category score, which could range from -3 to 3. A repeated measures ANOVA was conducted with category level (superordinate, subordinate, and unrelated) as repeated measures factors and age as a between subjects. A main effect of category level, $F(2, 138) = 19.04, p < .001, \eta^2 = .21$ was found, along with an effect of age, $F(2, 69) = 3.26, p = .045, \eta^2 = .086$. There was no significant interaction between age and category level. One sample t-tests (two-tailed) were therefore conducted to compare knowledge attributions at each categorical level to a chance value of 0 for each age group (see Table 2).

Categorical Level Age Group	Subordinate	Superordinate	Unrelated
1 (6-7)	M = .619 p = .056	M = .857 p = .007*	M = -.714 p = .015*
2 (8-9)	M = .1 p = .748	M = .9 p = .014*	M = -.4 p = .104
3 (adults)	M = .452 p = .232	M = 1.645 p < .001*	M = .032 p = .813

Table 2: Mean scores and p-values for one sample t-tests comparing knowledge attributions across each category level and age group to chance (* indicates $p < .05$)

Bonferroni-adjusted post hoc comparisons were also conducted to investigate how knowledge attributions at each categorical level differed from each other for each age group. For age group 1, mechanistic scores differed significantly between the superordinate and unrelated categories ($p = .011$), and between the subordinate and unrelated categories ($p = .023$), but not between the superordinate and subordinate categories ($p = 1.0$). For age group 2, mechanistic scores differed significantly between the superordinate and unrelated categories ($p = .025$) and marginally between the superordinate and subordinate categories ($p = .085$), but not between the subordinate and side categories ($p = .785$). For age group 3, mechanistic scores differed significantly between the superordinate and unrelated categories ($p < .001$), and between the superordinate and subordinate categories ($p = .045$), but not between the subordinate and unrelated categories ($p = .821$).

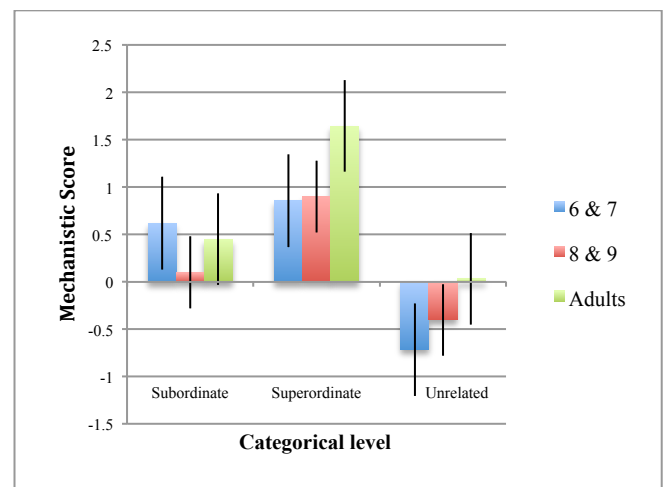


Figure 1: Study 1 mechanistic scores by categorical level and age group. Error bars indicate standard error.

Discussion

In support of our first hypothesis, each age group judged the mechanistic twin as more knowledgeable than the non-mechanistic twin about the superordinate level categories. In

support of our second hypothesis, no age group judged the mechanistic twin as more knowledgeable than the factual twin about the unrelated category. Our third hypothesis was mostly supported: no age group judged the mechanistic twin as more knowledgeable than the factual twin about the subordinate level category, although the youngest age group approached significance. Knowledge attributions at the superordinate and subordinate levels also did not differ significantly from each other for the youngest age groups, although scores at the superordinate level still exceeded those at the subordinate level for both age groups. In short, children and adults recognized that knowing how a category works implies broader knowledge within the same domain, but not greater knowledge about unrelated categories or, to an extent, the category itself.

Preferences for the unrelated category shifted across age groups, with the youngest children holding a non-mechanistic preference, the older children a comparatively weaker non-mechanistic preference, and adults possessing no preference. While children may have chosen the non-mechanistic twin strategically to balance knowledge attributions (i.e., from a fairness motivation to choose each twin some of the time), this might also reflect a legitimate preference.

There was also a developmental trend in subordinate level preferences. Younger children chose the mechanistic twin for the subordinate level almost as much as they did for the superordinate level. In contrast, older children and adults demonstrated a sizably weaker mechanistic preference at the subordinate level compared to the superordinate level. Thus, young children appear more optimistic about generalizing broad mechanistic knowledge to more specific instances of a category, an assumption that declines with age. This may reflect young children's lack of experience with specific mechanisms. As they begin to encounter them in school, children's intuitions about causal mechanisms become increasingly detailed and diverse, potentially weakening the perceived similarity between broad mechanistic knowledge and the mechanisms than obtain in specific kinds. Alternatively, this developmental trend might be due to children starting off with relatively less confidence that factual non-mechanistic information generalizes downward, and then an increasing awareness of this type of generalization with age.

In sum, the goal of this study was to investigate whether even young children appreciate that mechanistic knowledge generalizes beyond its immediate scope. We found evidence that children and adults selectively generalize mechanistic knowledge to immediate superordinate level categories, while still being aware of its limits within and across domains.

General Discussion

The current study provides evidence that children and adults attribute more knowledge about related kinds to individuals possessing mechanistic knowledge compared to factual non-mechanistic knowledge. In particular, we find evidence that

knowing how a basic level category works implies more superordinate category knowledge than simply knowing facts about it. This pattern suggests that a sense of shared mechanism among members of a kind guides children and adults' inferences about the knowledge others possess.

The study compared broadly mechanistic knowledge with factual non-mechanistic knowledge. One concern is that mechanistic knowledge was pitted against a particularly weak or idiosyncratic variety of non-mechanistic knowledge. To that end, the example pieces of knowledge varied widely and were chosen to give a broad coverage of factual non-mechanistic knowledge. Particular care was taken to avoid knowledge of surface features (such as color or size) that could be learned through mere observation. Instead, the facts concerned unobservable traits like history (e.g. where the first car engines were built) and constitution (e.g. what kinds of glass smartphones' screens are made out of). Also, if a particular component or topic was mentioned in a mechanistic example, it was also mentioned in the corresponding non-mechanistic example to minimize a mere bias for knowledge about internal parts. Importantly, the vignettes were explicitly labeled as example pieces of knowledge, and were meant to broadly indicative the kind of knowledge each twin possessed rather than specify it exactly. This was further implied when each twin was directly stipulated to have learned 'a lot' about the same basic level category.

A related concern is that mechanistic and non-mechanistic information is not truly dichotomous, either theoretically or cognitively. However, contemporary philosophers of science have offered broadly converging accounts (Bechtel, 2011; Craver & Darden, 2013). These accounts lay out several key features of mechanistic explanation in the sciences, including: a phenomenon being explained (e.g. how does it work?); a division of components, often functional, that underlie the phenomenon; a set of causal relations that obtain between the components, forming a bounded system; hierarchical organization of components via constitution, such that components can be unpacked into constituents and their interactions at a lower level. More colloquially, mechanistic explanations typically answer "how" and "why" questions and are compactly described as "how something works".

A selective preference for the mechanistic at the superordinate level shows that children and adults are capable of making distinct judgments about mechanistic knowledge when pitted against non-mechanistic factual knowledge, and that these judgments are reasonably bounded. Importantly, however, the distinction between mechanistic and non-mechanistic information is not always clear in everyday life. The two occur together in most contexts, and it is difficult to fully isolate a mechanism from relevant non-mechanistic specifications. For example, the size of a jet engine's parts directly impact how it works. However, the consistent mechanistic preferences at the superordinate level, despite an option to judge both knowledge types as equal, suggests that distinguishing

mechanistic from non-mechanistic knowledge is relatively easy and consequential when accompanied by a minimal level of supporting detail.

One limitation of these studies is that all stimulus items were artifacts, narrowing our findings to that domain. Artifacts were used for two reasons: first, mechanism is particularly salient in artifacts, often envisioned as the inner “clockworks” of objects (Dolnick, 2011); second, no colloquial phrase exists in English that signifies mechanistic knowledge about biological entities, in contrast to “how it works” for artifacts. Although, there is reason to suspect children generalize mechanistic knowledge across biological entities as well. Even infants believe the insides of living things have privileged causal powers (Newman et al, 2008), and by age six, children are able to make a variety of abstract inferences about an object’s internal features (Ahl & Keil, 2017). Taken together, these findings suggest some form of awareness of internal biological mechanisms by the early elementary school years. However, the biological domain may suffer from specific limitations; in particular, children’s, and even adults’, intuitions about plants are notably weak and delayed compared to their intuitions about animals (Stavy & Wax, 1989; Hatano & Inagaki, 1994). This might lead children to undervalue mechanistic knowledge about plants. Further research should test children’s mechanistic intuitions across a broad range of familiar biological entities, including animals, plants, and organs.

A final limitation is that, given children’s limited attention, the current studies did not feature multiple superordinate level categories for a single stimulus. As a result, it is unclear exactly how distant two categories within a domain need to be for mechanistic knowledge about one to no longer generalize to the other.

The scope of mechanistic knowledge is likely determined by the same intuitions that lead to it generalizing to related kinds, namely a sense of shared mechanism amongst those kinds. But what are these mechanistic intuitions like? How concrete and detailed are they? How do we acquire them and how do they change over time? The current studies do not address these questions, since the aim was to show that mechanistic knowledge generalizes, not what our representations of mechanisms are like. To some extent, these representations are idiosyncratic by nature, dependent on one’s concrete mechanistic knowledge and experiences with particular instances of a kind. However, given the relatively consistent pattern of responses in this study across all age groups, these representations may share fundamental features or structure in common. Future study of these common features could shed light on the nature of epistemic inference and conceptual cognition more broadly.

Conclusion

The current study examined how children and adults generalize mechanistic knowledge. The results suggest that even young children recognize that mechanistic knowledge about a basic level category implies greater knowledge

about its superordinate level category compared to factual non-mechanistic knowledge about the same basic level category. This provides one account of why intuitions about causal mechanism have such a strong influence on children and adults’ epistemic inferences. Impressively, children are able to make systematic inferences about mechanistic knowledge despite possessing little to none themselves, suggesting that abstract features of mechanism, rather than concrete mechanistic details, are chiefly responsible for mechanism’s influence on children and adults’ intuitions.

Acknowledgments

The authors thank the National Science Foundation who funded this research (proposal DRL 1561143 to F. Keil) as well as the members of the Yale Cognition and Development Lab for their helpful feedback.

References

- Ahl, R. E., & Keil, F. C. (2017). Diverse effects, complex causes: children use information about Machines' functional diversity to infer internal complexity. *Child Development, 88*(3), 828-845.
- Bechtel, W. (2011). Mechanism and biological explanation. *Philosophy of Science, 78*, 533- 557.
- Callanan, M. A., & Oakes, L. M. (1992). Preschoolers' questions and parents' explanations: Causal thinking in everyday activity. *Cognitive Development, 7*(2), 213-233.
- Chouinard, M. M., Harris, P. L., & Maratsos, M. P. (2007). Children's questions: A mechanism for cognitive development. *Monographs of the Society for Research in Child Development, i*-129.
- Corriveau, K. H., & Kurkul, K. E. (2014). “Why does rain fall?”: Children prefer to learn from an informant who uses noncircular explanations. *Child Development, 85*(5), 1827-1835.
- Craver, C. F., & Darden, L. (2013). In search of mechanisms: Discoveries across the life sciences. University of Chicago Press.
- Dolnick, E. (2011). The Clockwork Universe: saac Newto, Royal Society, and the Birth of the Modern World I. Harper Collins.
- Erickson, J. E., Keil, F. C., & Lockhart, K. L. (2010). Sensing the coherence of biology in contrast to psychology: Young children’s use of causal relations to distinguish two foundational domains. *Child Development, 81*(1), 390-409.
- 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.
- Hatano, G., & Inagaki, K. (1994). Young children's naive theory of biology. *Cognition, 50*(1-3), 171-188.
- Johnston, A. M., Sheskin, M., Johnson, S. G., & Keil, F. C. (2017). Preferences for Explanation Generality Develop Early in Biology But Not Physics. *Child Development*.
- Keil, F. C., Lockhart, K. L., & Schlegel, E. (2010). A bump on a bump? Emerging intuitions concerning the relative

- difficulty of the sciences. *Journal of Experimental Psychology: General*, 139(1), 1.
- Keil, F.C., Stein, C., Webb, L., Billings, V.D., & Rozenblit, L. (2008). Discerning the Division of Cognitive Labor: An Emerging Understanding of How Knowledge is Clustered in Other Minds. *Cognitive Science*, 32(2), 259-300.
- Kushnir, T., Vredenburg, C., & Schneider, L. A. (2013). "Who can help me fix this toy?" The distinction between causal knowledge and word knowledge guides preschoolers' selective requests for information. *Developmental Psychology*, 49(3), 446.
- Legare, C. H., & Lombrozo, T. (2014). Selective effects of explanation on learning during early childhood. *Journal of Experimental Child Psychology*, 126, 198-212.
- Lockhart, K. L., Chuey, A., Kerr, S., Keil, F. C. (under review). The Privileged Status of Knowing Mechanistic Information: An Early Epistemic Bias.
- Newman, G. E., Herrmann, P., Wynn, K., & Keil, F. C. (2008). Biases towards internal features in infants' reasoning about objects. *Cognition*, 107(2), 420-432.
- Sheskin, M. & Keil, F. C. (under review). TheChildLab.com: A video chat platform for developmental research.
- Stavy, R., & Wax, N. (1989). Children's conceptions of plants as living things. *Human Development*, 32(2), 88-94.
- Trouche, E., Chuey, A., Lockhart, K. L., & Keil, F. C. (2017). Why Teach How Things Work? Tracking the Evolution of Children's Intuitions about Complexity. *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (pp. 3368-3373). London, England, UK: Cognitive Science Society.
- Trouche, E., Chuey, A., Lockhart, K. L., & Keil, F. C. (under review). Why Teach How Things Work? Hidden Benefits of Exposure to Mechanistic Explanation
- 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.