Cognition 199 (2020) 104231

Contents lists available at ScienceDirect

Cognition

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

Children and adults selectively generalize mechanistic knowledge

Aaron Chuey[a,*], Kristi Lockhart[b], Mark Sheskin[b,c], Frank Keil[b]

[a] Stanford University, Psychology Department, United States of America
[b] Yale University, Psychology Department, United States of America
[c] Minerva Schools at KGI, School of Social Science, United States of America

A R T I C L E I N F O

Keywords:
Cognitive development
Knowledge
Mechanism
Epistemic development
Metacognition
Explanation

A B S T R A C T

A central component of evaluating others as sources of information involves estimating how much they know about different domains: one might be quite knowledgeable about a certain domain (e.g., clocks), but relatively ignorant about another (e.g., birds). Estimating one's domain knowledge often involves making inferences from specific instances or demonstrations, with some suggesting broader knowledge than others. For instance, an American who demonstrates knowledge of an unfamiliar country like Djibouti likely knows more about geography as a whole compared to an American who demonstrates knowledge of a more familiar country like Canada. The current studies investigate the extent to which one potentially salient kind of knowledge - mechanistic knowledge - signals greater domain knowledge as a whole. Across four developmental studies, we find that both adults and children as young as six think that those who possess mechanistic knowledge about a basic level artifact category (e.g., clocks) are more knowledgeable about its superordinate level category (e.g., machines) than those with factual non-mechanistic knowledge (Studies 1a and 2a). We also find an analogous, yet delayed pattern with biological categories (Studies 1b and 2b). Together, these studies demonstrate that even young children, who possess little mechanistic knowledge themselves, nevertheless have a sophisticated sense of how knowledge of mechanism generalizes across related categories.

1. Introduction

Information provided by other people can vary dramatically in its usefulness for a given problem. Selecting the most informative sources often relies on inferring each person's level of knowledge about the relevant domain. Frequently, the only basis for such inferences is whatever small amount of information each person has already revealed. The process of evaluating others' domain knowledge therefore depends on judgments about how different types of knowledge generalize. For example, if someone knows a lot about how tractors work, should we assume she is also knowledgeable about cars? Airplanes? Iguanas? Our intuitions about epistemic generalization not only help us identify which sources are knowledgeable, but also the limits of those sources' knowledge. A cellular biologist possesses extensive knowledge about how cells divide and grow, but may be just as ignorant about macroeconomics as any other non-economist. Thus, the way we infer how knowledge generalizes helps us gauge the knowledge of those around us, as well as predict the boundaries of others' expertise. Despite its large role in our epistemic inferences and whom we choose to defer to, little prior research has addressed what underlies intuitions about knowledge generalization. The current studies investigate one potentially early emerging influence – causal mechanism.

1.1. The importance of mechanism for generalization

A sense of shared causal mechanisms may play an important, early-emerging role in intuitions about how knowledge is clustered and how it generalizes. This role is suggested by demonstrations of the importance of mechanism in young children's inferential and epistemic strategies. By the preschool years, children possess a rudimentary understanding of causal mechanisms, distinguishing between relevant and irrelevant changes to a causal system (Buchanan & Sobel, 2011). This understanding pervades their epistemic considerations. 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), probably because children associate an ability to fix with mechanistic knowledge (Lockhart, Chuey, Kerr, & Keil, 2019).

Young children also utilize 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). Around the same time, children begin to cluster knowledge

* Corresponding author at: 450 Jane Stanford Way, Stanford, CA 94305, United States of America.
E-mail address: chuey@stanford.edu (A. Chuey).

https://doi.org/10.1016/j.cognition.2020.104231
Received 1 November 2018; Received in revised form 29 January 2020; Accepted 6 February 2020

domains in ways similar to how universities organize academic departments, such as separating biology and physics (Keil, Stein, Webb, Billings, & Rozenblit, 2008), implying a relationship between causal mechanism and knowledge organization. As children grow older, they 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 even their attitudes towards it.

Causal mechanism also influences 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; Frazier, Gelman, & Wellman, 2016), and those more general in scope (Johnston, Sheskin, Johnson, & Keil, 2018). 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, Harris, & Maratsos, 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). Young children also remember a larger number of causally relevant features when they explain phenomena, as opposed to merely reporting on them (Legare & Lombrozo, 2014; Walker, Lombrozo, Williams, Rafferty, & Gopnik, 2017), and they privilege "deep" properties in their own explanations, favoring more inductively powerful features of a system when engaging in explanation (Walker, Lombrozo, Legare, & Gopnik, 2014). All these studies suggest a link between mechanistic information and more encompassing forms of knowledge.

The appeal of mechanistic knowledge may arise from how it generalizes across related kinds, and young children seem to appreciate this distinctive property. Mechanism-focused instruction changes elementary school children's judgments about the complexity of an object more than factual non-mechanistic instruction (Trouche, Chuey, Lockhart, & Keil, 2017). Mechanism could cue children to certain properties that other forms of factual information do not. Further when children's complexity intuitions are influenced by mechanistic instruction, those changed intuitions also propagate to entities with related causal processes (Trouche, Chuey, Lockhart, & Keil, 2020). Children may therefore expect knowledge of mechanism-related properties to apply especially broadly. For example, if someone knows how trucks work, we might expect some of that knowledge to apply to tractors, because most vehicles work in broadly similar ways, utilizing analogous internal components.

The generalization of mechanistic knowledge does not preclude some degree of generalization for non-mechanistic knowledge as well. After all, almost all knowledge generalizes within a domain to some extent; knowing where trucks are built or how much they cost still implies some knowledge of vehicles more broadly. We argue here that, when making epistemic inferences, both adults and children tend to view mechanistic knowledge as a comparatively richer, more faithful cue to broader domain knowledge (i.e. knowledge about a superordinate level category) than other forms of knowledge.

1.2. The limits of generalizing mechanism

A nuanced understanding of how mechanistic knowledge generalizes includes not only an awareness that it generalizes, but also an awareness that it does not *always* generalize. For example, someone who knows how a car engine works probably understands broad automotive principles relevant to the operation of many kinds of vehicles (e.g., motorcycles, trucks, buses, etc.). In contrast, non-mechanistic knowledge about cars, like where car engines were invented, appears less applicable to other kinds of vehicles. Meanwhile, both mechanistic and non-mechanistic knowledge about cars seem equally inapplicable to an unrelated category like flowers. Importantly, mechanistic knowledge does not always generalize *more* than non-mechanistic knowledge. For example, knowing how car engines work entails some knowledge about how a particular subordinate level category, like race cars, work. Likewise, someone who possesses non-mechanistic knowledge about cars, like knowing where car engines were invented, could probably make a good guess about where race car engines were invented.

To summarize, although there are exceptions, mechanistic knowledge tends to generalize particularly well to the higher-level superordinate category. The broad principles applicable to car engine mechanisms often apply to many vehicles, while specific non-mechanistic facts about cars are generally limited to cars and not to other vehicles. In contrast, mechanistic and non-mechanistic knowledge both seem equally likely to generalize to specific lower-level subordinate categories; knowing facts about cars or how they work implies further knowledge about specific kinds of cars. Finally, mechanistic and non-mechanistic knowledge are both equally unlikely to generalize to unrelated categories; any kind of car-related knowledge is just as unlikely to apply to flowers.

1.3. Mechanism across domains

Although one might possess mechanistic knowledge across many domains, mechanism is most salient in domains with a rich hierarchical structure consisting of functional components with underlying causal relations, such as artifacts and biological systems. However, despite pervasive interest in biological mechanisms by philosophers of science (see Bechtel, 2011; Kaiser & Krickel, 2017), children and adults may have weaker intuitions for how mechanistic knowledge about biological categories generalizes. Mechanism is particularly apparent in artifacts and often envisioned as the inner "clockworks" of objects (Dolnick, 2011). In contrast, although animals and plants have rich hierarchical structures of nested functions and supporting causal processes, their mechanistic subcomponents (organs, cells, etc.) are more obscure, at least to children. Mechanistic similarities among biological kinds might therefore be more difficult to discern than for artifact kinds, which have more visually accessible subcomponents and more immediately obvious mechanistic similarities.

Although biological mechanisms may be more enigmatic to children, they might nonetheless generalize mechanistic knowledge to related biological categories. Even infants believe the insides of living things have privileged causal powers (Newman, Herrmann, Wynn, & Keil, 2008; Taborda-Osorio & Cheries, 2017), and by age six, children are able to make a variety of abstract inferences about a novel object's internal features (Ahl & Keil, 2017). If children are able to apply their understanding of the insides of objects to the insides of biological entities, then they may possess some form of awareness of internal biological mechanisms by the early elementary school years. Thus, in the current studies, we compare identical experiments with artifactual and biological stimuli to investigate whether children and adults generalize mechanistic knowledge differently across these domains. Because children's, and even adults', intuitions about plants are notably weak and delayed compared to animals (Hatano & Inagaki, 1994; Stavy & Wax, 1989), we used familiar animal categories for the biological domain.

1.4. The current studies

In the current studies, we investigate when individuals across three age groups — six- and seven-year-olds, eight- and nine-year-olds, and adults — think a person who possesses mechanistic knowledge about a category knows more about another category than someone possessing factual non-mechanistic knowledge. We examine participants' epistemic judgments about superordinate, subordinate, and unrelated categories. If participants think the person with mechanistic knowledge knows more about all three categories, such intuitions would suggest an appreciation for the importance of mechanism *without* an appreciation of its limits. Alternatively, participants might limit their generalization of mechanistic knowledge. Participants might choose the person who knows mechanism as more knowledgeable for both of the related domains, superordinate and subordinate, but not for the unrelated domain; or they might choose the person who knows about mechanism as knowing more about just the superordinate domain. This final possibility, choosing the mechanistic knower for the superordinate category only, would imply that their knowledge generalizes more broadly across the domain, without being more knowledgeable about subcategories of the basic level category itself.

We included multiple age groups to assess whether generalizing mechanistic knowledge follows an early or late developmental trajectory. While previous studies have shown that children as young as six extensively associate mechanistic knowledge with greater expertise (Lockhart, Chuey, Kerr, & Keil, 2019), the breadth of children's knowledge attributions in these cases is unclear. In addition, six- and seven-year-olds are at the earliest stages of education, and usually receive no instruction about specific mechanisms until later in elementary school (see Next Generation Science Standards: Lead States, 2013). Thus, a systematic ability to generalize mechanistic knowledge at these younger ages would indicate that mature intuitions do not merely arise from possessing detailed mechanistic knowledge. Rather, they might arise from observations of those who do have that knowledge, or even from more abstract ideas about how knowledge is acquired or clustered.

Eight- and nine-year-olds were included to shed light on the developmental trajectory of these intuitions. Previous work has shown that children's category-based induction continues to develop significantly into the elementary years. For example, children younger than eight have difficulty understanding how many subcategories two categories share (López, Gelman, Gutheil, & Smith, 1992). Children's inferences about knowledge generalizability may follow a parallel development.

While adults were included as a comparison group against children's judgments, they also serve as a group of interest in their own right. No prior work has investigated how adults generalize mechanistic knowledge within versus across domains, making the nature of the mature intuition itself an important empirical question.

2. Study 1a

Participants in the first study were young children (six- and seven-year-olds), older children (eight- and nine-year-olds), and adults. All participants heard a story about two twins, one possessing mechanistic knowledge about a basic level artifact category (e.g., clocks) and the other possessing factual non-mechanistic knowledge about that category. We then asked participants 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). We used twins to imply that both characters were otherwise identical.

The superordinate category was our key measure. Our primary hypothesis was that participants would judge the twin possessing mechanistic knowledge as more knowledgeable about the superordinate category than the twin possessing factual non-mechanistic knowledge.

The subordinate category was included to assess the scope of the generalization inferences. Participants might think the mechanistic twin knows more about the subordinate category. On the other hand, participants might think that the two twins are likely to have approximately equal knowledge about the subordinate category. The subordinate category therefore clarifies whether participants think mechanistic knowledge generalizes more than factual knowledge to a related subcategory, or whether it selectively generalizes up to the superordinate category.

We included the unrelated basic level category in case participants do not distinguish between the superordinate and subordinate categories. If children choose the mechanistic twin across all three domains (superordinate, subordinate, and unrelated), this would suggest they see the mechanistic knower as simply more knowledgeable about anything.

We predicted that older children and adults would infer that mechanistic knowledge about artifacts generalizes selectively, to the superordinate level only. In contrast, younger children have demonstrated a more unconstrained bias for mechanistic information in previous tasks. In one such task, Lockhart, Chuey, Kerr, and Keil (2019) presented children with an individual who possessed mechanistic knowledge about an object and another who possessed knowledge relevant to marketing it. They were then asked who they would choose to help them either fix or sell the object. While young children chose the mechanistic knower significantly more to help them fix the object than sell it, they still preferred the mechanistic knower for both goals. Young children might generalize mechanistic knowledge in a similarly unconstrained manner. Thus, we predicted younger children would generalize mechanistic knowledge about artifacts more broadly, to both the superordinate and subordinate levels.

2.1. Method

2.1.1. Participants

We conducted power analyses for both families of tests planned for this study (repeated measures ANOVAs powered for a within-between interaction and two-tailed one-sample *t*-tests). For the repeated measures ANOVA (3 groups, 3 measurements), to find a significant ($\alpha = 0.05$) within-between interaction with a medium effect size ($f = 0.25$) and medium power ($1 - \beta = 0.8$), only 36 total participants (12 per age group) were necessary. For the one-sample *t*-tests (two-tailed), we found that to find a significant ($\alpha = 0.05$) medium effect ($d = 0.5$) with medium power ($1 - \beta = 0.8$), 34 participants per age group were necessary. We used the higher estimate and added an extra participant to each age group. A post hoc power analysis is included in *Results* to measure whether adequate power was achieved.

35 six- and seven-year-olds (M_{age}: 6 years 10 months, range: 73–92 months, 23 male) and 35 eight- and nine-year-olds (M_{age}: 8 years 10 months, range: 97–119 months, 17 male) participated in the experiment; one child was excluded, with replacement, due to experimenter error. All children participated via TheChildLab.com online platform (Sheskin & Keil, 2020). 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, such as following a ball through a tube, which were established as easy for most child participants in previous research. 49 adults participated in the experiment via Amazon Mechanical Turk for \$0.50 payment; 41 adults passed standard attention checks. The first 35 (M_{age}: 35 years, Range: 22–56 years, 24 male) were included in the final sample to match the size of each child sample, although our conclusions are identical if all adults are included.

2.1.2. Materials

We used three stimulus categories, with each matched to a different superordinate category, subordinate category, and unrelated basic level category. The three sets (stimulus/superordinate/subordinate/unrelated) were: 1) clocks, machines, grandfather clocks, tulips; 2) cars, (wheeled) vehicles, race cars, sharks; 3) smartphones, electronics, iPhones, tigers. Each category was presented with an image depicting the category, consisting of six category exemplars in a white square (to emphasize kind rather than token, see Fig. 1). These categories were chosen because they represented a broad sample of artifacts familiar to most children. Each item had corresponding mechanistic and non-mechanistic knowledge vignettes (see Table 1), which exemplified each twin's knowledge. When possible, the non-mechanistic examples referenced the same internal component or topic as the corresponding mechanistic example. The non-mechanistic examples were also designed to reference information that could not be known via

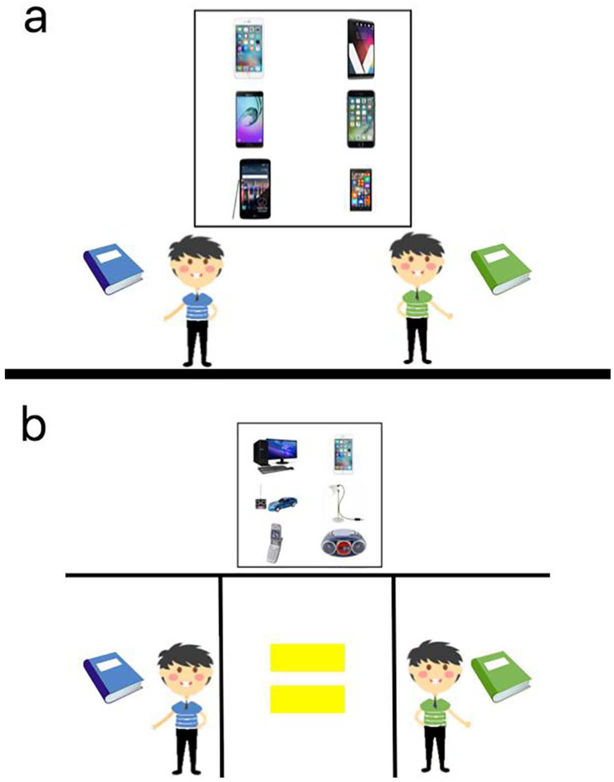

Fig. 1. Example slides from Study 1. (a) Each twin reads a book about a category ("smartphones", with 6 exemplars in a 2 × 3 grid). (b) Participants are asked which twin (blue or green) knows more about the superordinate level category (electronics, with 6 exemplars in a 2 × 3 grid, including 1 smartphone), or if they both knew the same amount (yellow equals sign). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

observation alone, such as an object's history or constitution, as well as quantitative information.

2.1.3. Design

Each question was focused on which one of two cartoon children knew more about a kind. Each pair was introduced as twins, and looked nearly identical, except that one twin wore blue clothes and the other wore green clothes. The twins were referred to by the color of their clothes, as "Blue" or "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, unrelated) were presented in a consistent order for a given participant across all three 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 min for children and 5 min for adults.

2.1.4. Procedure

At the start of the activity, we presented children with a training example that introduced the concept of a yellow equal sign, which would be used in the activity. Participants then learned 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. Instead of this training, adults received the following written instructions: "*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*."

We then introduced participants to both twins and told them that each twin read a book, with colors matching their corresponding twin; 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. We told participants that one twin learned about how the category works (the mechanistic twin), and the other twin learned facts about the category (the non-mechanistic twin). We also provided two examples of each twin's mechanistic or non-mechanistic knowledge (see Table 1).

Next, we presented participants with a test category and asked them which twin knew more about the category, or whether 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?*" We asked children about the test categories in sequence, each presented on a different slide. We presented adults with all three test categories in sequence on the same page. Participants completed this procedure for each stimulus category.

2.2. Results

In our analyses, choosing the mechanistic twin as knowing more was coded as 1, choosing the non-mechanistic twin as knowing more was coded as -1, and choosing both twins as knowing the same amount was coded 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, each of which could range from -3 to 3.

Repeated measures ANOVA with mechanistic scores at each category level (superordinate, subordinate, and unrelated) as within subjects factors and age as a between subjects factor found a main effect of category level, $F(2, 204) = 22.71$, $p < .001$, $\eta 2 = 0.13$. There was also a significant effect of age, $F(2, 102) = 5.24$, $p = .007$, $\eta 2 = 0.09$, but the interaction was not significant, $F(4, 204) = 1.86$, $p = .12$.

Separate one sample *t*-tests (two-tailed) were conducted to compare mechanistic scores at each categorical level to a chance score of 0 for each age group. At the superordinate level, the mechanistic scores of all

Table 1
Study 1a knowledge examples.

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 cars 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.

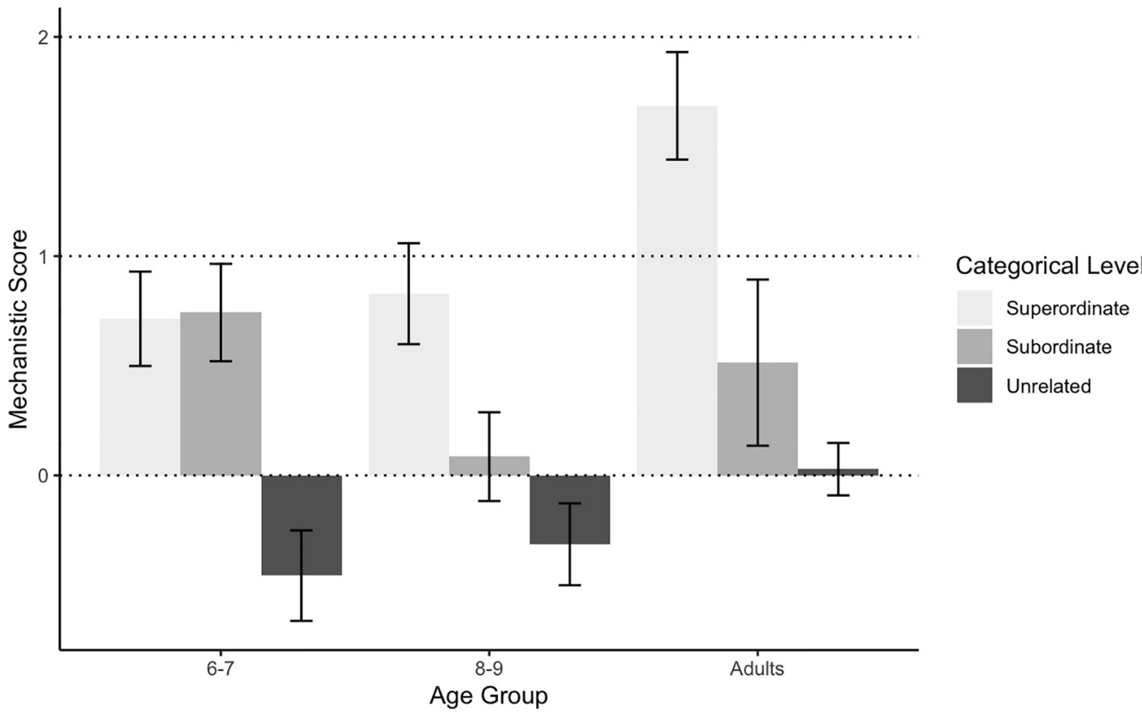

Fig. 2. Study 1a mechanistic scores by categorical level and age. Error bars indicate standard error.

age groups were significantly above chance: six- and seven-year-olds, ($M = 0.71$, $SD = 1.27$) $t(34) = 3.32$, $p = .002$; eight- and nine-year-olds, ($M = 0.83$, $SD = 1.36$), $t(34) = 3.60$, $p < .001$; adults, ($M = 1.69$, $SD = 1.45$), $t(34) = 6.88$, $p < .001$. At the subordinate level, the mechanistic scores of six- and seven-year-olds were significantly greater than chance ($M = 0.74$, $SD = 1.31$), $t(34) = 3.35$, $p = .002$, while eight- and nine-year-olds, ($M = 0.09$, $SD = 1.20$), $t(34) = 0.42$, $p = .41$, and adults ($M = 0.51$, $SD = 2.24$), $t(34) = 1.36$, $p = .18$ did not differ significantly from chance. At the unrelated level, the mechanistic scores of six- and seven-year-olds ($M = -0.46$, $SD = 1.22$) were significantly below chance, $t(34) = -2.22$, $p = .034$, while eight- and nine-year-olds, ($M = -0.31$, $SD = 1.12$), $t(34) = -1.68$, $p = .1$, and adults, ($M = 0.03$, $SD = 0.71$), $t(34) = 0.24$, $p = .81$, did not differ significantly from chance.

To test whether we achieved adequate power given our sample and effect sizes, we conducted a post hoc power analysis (ANOVA: repeated measures, within-between interaction) based on the effect size of the interaction between categorical level and age group ($f = 0.19$). Ample power ($1 - \beta = 0.98$) was achieved given the size of our sample and effect of the interaction (Fig. 2).

2.3. Discussion

All age groups judged that the mechanistic twin was more knowledgeable than the non-mechanistic twin at the superordinate level. Importantly, mechanistic scores universally were higher at the superordinate level than the unrelated level, which did not differ from chance (older children and adults) or were significantly less than chance (younger children). However, judgments at the subordinate level varied. Older children and adults' mechanistic scores at the subordinate level did not differ from chance, demonstrating a selective sense of generalization. Conversely, younger children's mechanistic scores at the subordinate level were significantly greater than chance, suggesting a more unconstrained sense of generalization overall.

In sum, by the elementary school years, children generalize mechanistic knowledge about a basic level artifact category to its superordinate level category. In effect, knowing how a kind of artifact works, compared to knowing facts about it, licenses greater knowledge about other categories of that kind as well. Importantly, even the youngest age group did not extend this judgment to the unrelated category. Younger children, but not older children and adults, judged the mechanistic twin as more knowledgeable at the subordinate level, suggesting that a refined sense for how mechanistic knowledge about artifacts generalizes continues to develop throughout the elementary school years.

3. Study 1b

While Study 1a provided evidence for our primary hypothesis with respect to artifact categories, it remains unclear whether children and adults selectively generalize mechanistic knowledge within other domains as well. Thus, Study 1b investigates children and adults' generalization intuitions surrounding biological categories. Similar to Study 1a, we predicted that adults would selectively generalize mechanistic knowledge about an animal category to its superordinate level category. However, because knowledge of mechanism is more obscure and less perceptually salient for biological entities, we predicted that only older children would selectively generalize mechanistic knowledge to the superordinate level and that younger children would show no significant preference for mechanistic knowledge at the superordinate level.

3.1. Method

3.1.1. Participants

35 six- and seven-year-olds (M_{age}: 7 years 0 months, range: 73–95 months, 20 male) and 35 eight- and nine-year-olds (M_{age}: 9 years 1 month, range: 97–118 months, 15 male) participated in the experiment; no children were excluded. They participated via TheChildLab.com online platform (Sheskin & Keil, 2020). 40 adults participated in the experiment via Amazon Mechanical Turk for $0.50 payment; 38 adults passed standard attention checks; the first 35 (M_{age}: 37 years, Range: 21–69 years, 22 male) were included in the final sample to match the size of each child sample, but our conclusions are identical if

Table 2
Study 1b knowledge examples.

Stimulus	Mechanistic knowledge	Non-mechanistic knowledge
Sparrows	For example, she learned how Sparrows' stomachs break down hard seeds and insects. As another example, she learned how sparrows' voice boxes vibrate to make chirping sounds.	For example, she learned how long it normally takes sparrows to search for food. As another example, she learned at what time of day Sparrows chirp the most.
Sharks	For example, he learned how sharks are able to constantly grow new teeth. As another example, he learned how sharks' gills work so they can breathe underwater.	For example, he learned how large shark teeth get. As another example, he learned how many gills sharks have on each side of their body.
Tigers	For example, he learned how tigers' eyes are able to see in the dark. As another example, he learned how tigers' stomachs absorb nutrients from meat.	For example, he learned the different kinds of colors tigers' eyes can be. As another example, he learned how much meat can fit in tigers' stomachs.

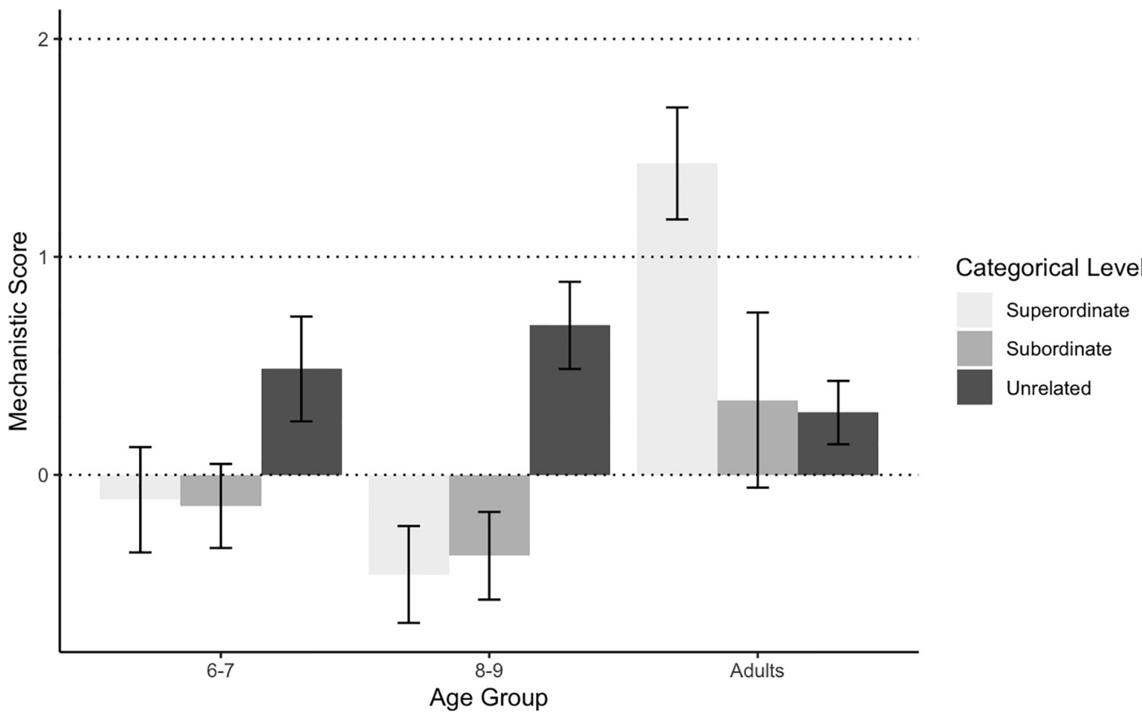

Fig. 3. Study 1b mechanistic scores by categorical level and age. Error bars indicate standard error.

all adults are included.

3.1.2. Materials

We used three stimulus categories, with each matched to a different superordinate level category, subordinate level category, and unrelated basic level category. The three sets (stimulus/superordinate/subordinate/unrelated) were: 1) sparrows, birds, java sparrows, clocks; 2) sharks, fish, bull sharks, cars; 3) tigers, carnivores, Siberian tigers, smartphones. Each category was presented with an image depicting the category, consisting of six category exemplars in a white square. These categories were chosen because they represented a broad sample of animals familiar to most children. Each item had corresponding mechanistic and non-mechanistic knowledge vignettes (see Table 2), which exemplified each twin's knowledge. When possible, the non-mechanistic examples referenced the same body part or function as the corresponding mechanistic example. The non-mechanistic examples were also designed to reference information that could not be known via observation alone, such as an animal's abilities and tendencies, or quantitative information about its body parts.

3.1.3. Design

As in Study 1a, knowledge type (mechanistic or non-mechanistic) and which categorical level was presented first (superordinate or subordinate) were independently counterbalanced across participants while the order of the stimulus items was randomized across participants. The study took approximately the same amount of time as Study 1a (8 min for children and 5 min for adults).

3.1.4. Procedure

The experimental procedure was identical to that of Study 1a, except using biological stimulus categories.

3.2. Results

The same coding scheme as Study 1a was used, yielding a superordinate level score, a subordinate level score, and an unrelated category score, aggregated across stimulus items, that could range from -3 to 3. Repeated measures ANOVA with mechanistic scores at each category level (superordinate, subordinate, and unrelated) as within subjects factors and age as a between subjects factor found a main effect of category level, $F(2, 204) = 3.81$, $p = .024$, $\eta 2 = 0.02$. There was also a significant effect of age, $F(2, 102) = 7.91$, $p < .001$, $\eta 2 = 0.13$. The interaction was not significant, $F(4, 204) = 6.06$, $p = .069$.

Separate one sample *t*-tests (two-tailed) were conducted to compare mechanistic scores at each categorical level to a chance score of 0 for each age group. At the superordinate level, only adults' mechanistic scores ($M = 1.43$, $SD = 1.52$) were significantly above chance, $t(34) = 5.56$, $p < .001$. The scores of six- and seven-year-olds ($M = -0.114$, $SD = 1.43$) did not differ from chance, $t(34) = -1.43$, $p = .64$; and the scores of eight- and nine-year-olds ($M = -0.46$, $SD = 1.31$) were significantly below chance, $t(34) = -2.06$, $p = .047$. At the subordinate level, none of the mechanistic scores for the three age groups differed significantly from chance: six- and seven-year-olds ($M = -0.14$, $SD = 1.14$), $t(34) = -0.74$, $p = .46$; eight- and nine-year-olds ($M = -0.37$, $SD = 1.19$), $t(34) = -1.85$, $p = .074$; adults ($M = 0.34$, $SD = 2.38$), $t(34) = 0.85$, $p = .4$. At the unrelated level, the mechanistic scores of eight- and nine-year-olds ($M = 0.69$, $SD = 1.18$), $t(34) = 3.43$, $p = .002$ were significantly greater than chance. There was also a tendency for the mechanistic scores of six- and seven-year-olds ($M = 0.49$, $SD = 1.42$) and adults ($M = 0.29$, $SD = 0.86$) to be greater than chance as well, $t(34) = 2.02$, $p = .051$, and $t(34) = 1.97$, $p = .058$, respectively (Fig. 3).

3.3. Discussion

As in Study 1a, adults selectively generalized mechanistic knowledge about a basic level animal category to its superordinate level category. Children, however, differed in their judgments about how biological mechanistic knowledge generalizes. Older children thought the twin with non-mechanistic knowledge knew more about the superordinate level category. In contrast, they judged the mechanistic twin to know significantly more about the unrelated artifact category. Although younger children's judgments were generally at chance, they also judged the mechanistic twin to know marginally more about the unrelated artifact category, a tendency adults displayed as well. To summarize, adults generalize mechanistic knowledge about biological categories and artifact categories similarly, selectively generalizing both to the superordinate level. However, children judged the twin possessing mechanistic biological knowledge as only more knowledgeable about the unrelated artifact category. Additionally, older children selectively generalized non-mechanistic knowledge about biological categories to the superordinate level.

Children's conflicting performance in Studies 1a and 1b could reflect disparate expectations about how knowledge generalizes to artifact and biological categories. Perhaps children more strongly associate knowledge of mechanism with artifact knowledge, and knowledge of “facts” with biological knowledge. This seems plausible given the abundance of fun fact-oriented shows about animals that are intended for child audiences. Alternatively, older children's performance in Study 1b could have been exclusively driven by a strong association between mechanistic and artifact knowledge. Children might generalize any kind of mechanistic knowledge (including biology) to knowledge of artifacts more broadly, causing them to attribute knowledge about an unrelated artifact category to the mechanistic twin in Study 1b. In the child's mind, someone who knows all about the mechanics of gills, eyes and voice boxes might be an especially “mechanistically minded” person who is more likely to know about other overtly mechanical objects, such as clocks, cars and smartphones. In turn, children might map the non-mechanistic twin to the remaining biological categories by default in order to balance knowledge attributions across both twins.

To differentiate between these alternatives, Studies 2a and 2b replicate the previous studies using a simpler experimental paradigm that only tests participants' judgments about the superordinate level category with no contrasting categories from other domains. This allows us to distinguish whether children generalize knowledge differently for artifact and biological categories, or if the previously observed difference was simply driven by an association between knowledge of mechanism and artifacts. Additionally, the previous studies included the option to respond that the twins knew either the same or different amounts, which may have posed additional cognitive loads unrelated to the problem of how knowledge generalizes. Therefore, Studies 2a and 2b employ a binary choice measure with no middle option. If children's disparate performance in Studies 1a or 1b was driven by a genuine difference in knowledge attributions for artifacts and biological categories, they should still generalize mechanistic knowledge to the superordinate level for artifact categories and non-mechanistic knowledge to the superordinate level for biological categories when presented in isolation. However, if their answers in 1b were caused by a strong association between mechanistic and artifact knowledge, that effect should disappear, or even reverse, without the presence of artifact categories.

4. Study 2a

Study 1a demonstrated that children as young as six selectively generalize mechanistic knowledge from a basic level artifact category to its superordinate level category. Study 2a replicates the core result of Study 1a, but using a simpler method. Specifically, the twins in Study 2a learned about a basic level category, as in Study 1a, but we asked about only one target category at the superordinate level, not subordinate or unrelated, and provided a dichotomous choice with no equal option. Asking about a single categorical level eliminates the possibility that our original effect was driven by the presence of contrasting categorical levels, in addition to reducing attentional demands and the overall time of the task. A dichotomous choice also decreases the complexity of the task by providing a simpler preference measure.

4.1. Method

4.1.1. Participants

35 six- and seven-year-olds (M_{age}: 6 years 10 months, range: 72–94 months, 19 male) and 35 eight- and nine-year-olds (M_{age}: 8 years 10 months, range: 98–119 months, 17 male) participated in the experiment; one child was excluded, with replacement, due to experimenter error. They participated via TheChildLab.com online platform (Sheskin & Keil, 2020). 40 adults participated in the experiment via Amazon Mechanical Turk for $0.50 payment; 38 adults passed standard attention checks; the first 35 (M_{age}: 34 years, Range: 21–66 years, 19 male) were included in the final sample to match the size of each child sample, but our conclusions are identical if all adults are included.

4.1.2. Materials

The same twins, test categories (see Fig. 1), and knowledge examples (see Table 1) from Study 1 were used, with only the superordinate level categories being tested.

4.1.3. Design

Study 2a was similar to 1a and used the same stimuli, except we only tested one category (superordinate level) for each item. Additionally, we eliminated the neutral answer choice (i.e. the yellow equals sign indicating that both twins knew the same amount), which served to simplify the instructions and answer choices. The blue twin was always presented first for each stimulus token, with knowledge type counterbalanced. The order of presentation for the stimulus tokens was randomized. The study took approximately 4 min for children and 3 min for adults.

4.1.4. Procedure

Because Study 2a was a binary choice paradigm, we trained children to answer “blue” for the blue twin and “green” for the green twin, but did not need to introduce or use the yellow equals sign. We instructed adults: “*In this survey, you are going to hear about pairs of twins who both learn a lot things about a topic, but the things they learn about it are different. You will then be asked who you think knows more about the kind of thing. For example, imagine two twins each learn some different things about the refrigerator in their kitchen. You will then be asked who you think knows more about refrigerators.*”

We then presented participants the first pair of twins and a basic level category with accompanying exemplars. We told participants that both twins really knew nothing about the category (e.g., clocks), so a parent decided to teach them some things about it, but the parent taught different things to each twin. One twin learned how the category works, while the other learned facts about it. Both descriptions of the twins' knowledge were accompanied by two example pieces of their knowledge (see Table 1). Afterwards, we told participants that both twins now knew some things about the category. We then asked which twin knew more about its superordinate level category. We repeated this procedure for the other two items.

4.2. Results

In our analyses, choosing the mechanistic twin was coded as 1 and the non-mechanistic twin as 0, yielding a mechanistic score aggregated across stimuli that ranged from 0 to 3. An ANOVA with mechanistic score as the dependent variable and age group as the fixed factor revealed a significant effect of age group $F(2, 102) = 22.89$, $p < .001$, $\eta 2 = 0.31$. Post hoc comparisons demonstrated that the effect was driven by mechanistic scores increasing linearly with age. Eight- and nine-year-olds had significantly higher mechanistic scores ($M = 2.29$, $SD = 0.67$) than six- and seven-year-olds ($M = 1.83$, $SD = 0.71$), Bonferroni $p = .008$. Adults ($M = 2.83$ $SD = 0.45$) had significantly higher mechanistic scores than both six- and seven-year-olds and eight- and nine-year-olds, Bonferroni $p = .001$ and $p < .001$, respectively. One sample t-tests (two-tailed) were conducted comparing mechanistic scores to a chance value of 1.5 for each age group. The mechanistic scores of all age groups differed significantly from chance: six- and seven-year-olds, $t(34) = 2.75$, $p = .009$; eight- and nine-year-olds $t(34) = 6.97$, $p < .001$; adults $t(34) = 17.36$, $p < .001$.

4.3. Discussion

Replicating the results of Study 1a, all three age groups judged that the twin who possessed mechanistic knowledge about a basic level artifact category knew comparatively more about its superordinate category than someone possessing non-mechanistic factual knowledge. There was also a tendency to generalize mechanistic knowledge more strongly with age, indicating that exposure to mechanistic details, particularly through formal education, may continue to shape children's intuitions into adulthood (see Fig. 4).

5. Study 2b

Study 2a replicated the results of Study 1a with a simpler paradigm. It also found an increasing tendency to generalize mechanistic knowledge to the superordinate level with age, indicating the simpler task may be a more sensitive measure of children's generalization intuitions. Study 2b uses the same method to investigate children and adults' generalization intuitions about biological categories. If children continue to generalize non-mechanistic biological knowledge to the superordinate level in the absence of artifact categories, this would corroborate the results of Study 1b, suggesting children indeed have conflicting intuitions about how mechanistic knowledge generalizes across artifact and biological categories. However, if children do not display this tendency, then older children's performance in Study 1b was likely caused by their strong association between mechanistic and artifact knowledge rather than a genuine tendency to generalize non-mechanistic biological knowledge to the superordinate level.

5.1. Method

5.1.1. Participants

35 six- and seven-year-olds (M_{age}: 7 years 0 months, range: 73–93 months, 17 male) and 35 eight- and nine-year-olds (M_{age}: 9 years 1 month, range: 97–119 months, 25 male) participated in the experiment; one child was excluded, with replacement, due to technical difficulties. They participated via TheChildLab.com online platform (Sheskin & Keil, 2020). 40 adults participated in the experiment via Amazon Mechanical Turk for $0.50 payment; 37 adults passed standard attention checks; the first 35 (M_{age}: 38 years, Range: 20–73 years, 23 male) were included in the final sample to match the size of each child sample, but our conclusions are identical if all adults are included.

5.1.2. Materials

The same twins, test categories, and knowledge examples from Study 1b were used (see Table 2), with only the superordinate level categories being tested.

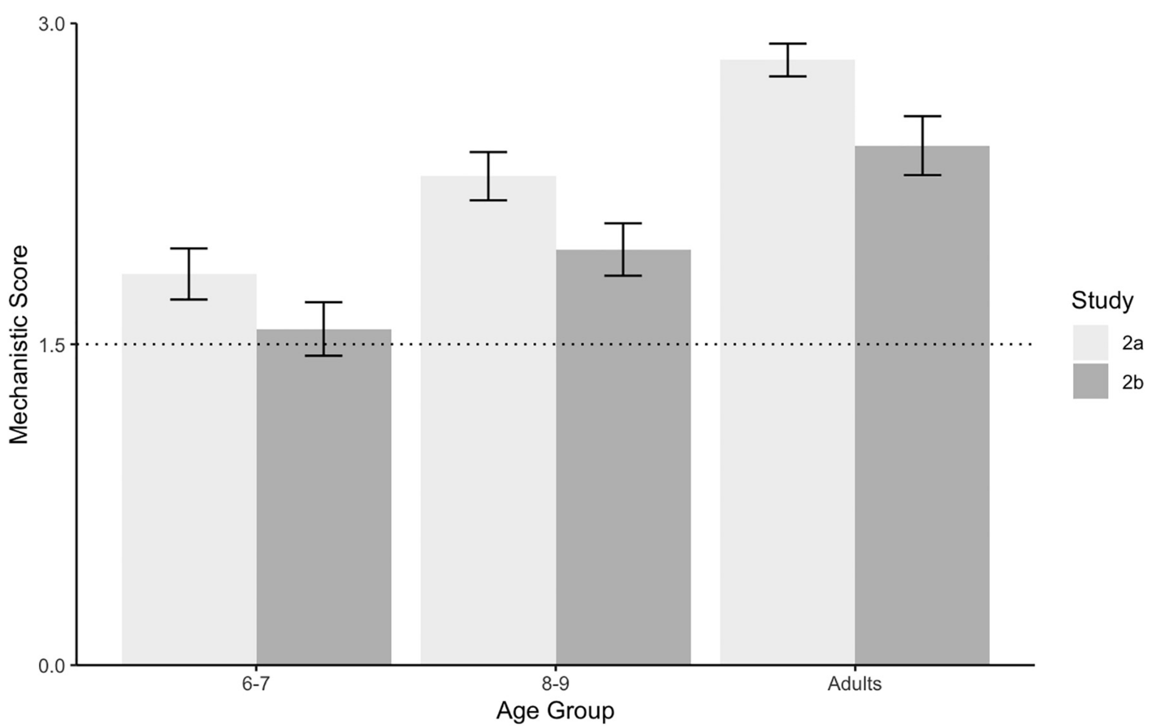

Fig. 4. Study 2a and 2b mechanistic scores by age. Error bars indicate standard error.

5.1.3. Design

Study 2b utilized the same design as Study 2a.

5.1.4. Procedure

The procedure was identical to 2a, except using biological stimulus items.

5.2. Results

In our analyses, choosing the mechanistic twin was coded as 1 and the non-mechanistic twin as 0, yielding a mechanistic score aggregated across stimuli that ranged from 0 to 3. An ANOVA with mechanistic score as the dependent variable and age group as the fixed factor revealed a significant effect of age group $F(2, 102) = 11.17$, $p < .001$, $\eta 2 = 0.18$. Post hoc comparisons demonstrated that the effect was driven by mechanistic scores increasing linearly with age; adults had significantly higher mechanistic scores ($M = 2.43$ $SD = 0.82$) than six- and seven-year-olds ($M = 1.57$, $SD = 0.74$), Bonferroni $p < .001$, and eight- and nine-year-olds ($M = 1.94$, $SD = 0.73$), Bonferroni $p = .026$. The difference between older and younger children's mechanistic scores was not significant, Bonferroni $p = .13$.

One sample t-tests (two-tailed) were conducted comparing mechanistic scores to a chance value of 1.5 for each age group. The mechanistic scores of eight- and nine-year-olds, $t(34) = 3.61$, $p < .001$, and adults, $t(34) = 17.36$, $p < .001$ differed significantly from chance. The mechanistic scores of six- and seven-year-olds did not differ from chance: $t(34) = 0.57$, $p = .57$.

A separate ANOVA with mechanistic score as the dependent variable and age group and study as fixed factors compared knowledge attributions in Study 2a and 2b. There was a significant effect of age, $F(2, 204) = 31.5$, $p < .001$, $\eta 2 = 0.23$, and study, $F(1, 204) = 12.13$, $p < .001$, $\eta 2 = 0.043$. The interaction was not significant $F(2, 204) = 0.19$, $p = .83$. Bonferroni adjusted post hoc comparisons showed that the effect of study was driven by significantly higher mechanistic scores in Study 2a ($M = 2.31$, $SD = 0.74$) compared to 2b ($M = 1.98$, $SD = 0.83$), Bonferroni $p = .002$.

5.3. Discussion

When assessed via a simpler, binary-choice task using the same stimuli as Study 1b but with no contrasting artifact categories, older children now show a significant preference for the mechanistic twin at the superordinate level. Therefore, the presence of artifact categories appears to have obscured older children's preferences in Study 1b. However, younger children still performed at chance levels. Mechanistic scores were also uniformly higher in Study 2a than 2b across all age groups, suggesting that mechanistic knowledge may be more strongly associated with artifact knowledge across the lifespan. Thus, the mature intuition itself appears weaker for biological categories, and develops later.

6. General discussion

The current studies provide evidence that children and adults think knowledge about mechanism generalizes more broadly than knowledge about non-mechanistic facts in the artifact and biological domains. In particular, we find that children and adults attribute more superordinate category knowledge to someone who knows how a basic level artifact category works, compared to someone who knows facts about that category. Importantly, this result does not hold for unrelated categories, indicating that children and adults do not merely assume knowledge of mechanism makes someone smarter or more knowledgeable in general.

For biological categories, children appear to more slowly develop a mature sense of how mechanistic knowledge generalizes to the superordinate level category. Adults selectively judged that mechanistic knowledge about a basic level biological category implies greater knowledge about its superordinate level category in both Studies 1b and 2b. Older children made the opposite judgment in the presence of contrasting artifact categories (Study 1b), but shared adults' intuitions in a simpler, binary-choice task (Study 2b). Meanwhile, younger children showed no significant preference in either task. Children's developing intuitions about mechanistic knowledge in the biological domain during the elementary school years may support more refined generalization judgments during the same period. By contrast, children's intuitions about mechanistic knowledge in the artifact domain appear to be sufficiently well developed at the beginning of formal education to serve as a basis for epistemic generalization.

6.1. Generalizing mechanistic knowledge

Children as young as six have systemic intuitions about how mechanistic knowledge generalizes, despite knowing very little about actual mechanisms themselves. These kinds of intuitions may therefore not arise from an early grasp of specific mechanisms. Instead, children might develop their intuitions by observing others who possess mechanistic knowledge and by making inductions on the basis of the other kinds of knowledge those individuals also demonstrate. For example, if a child observes their grandfather fix a car one day, a boat the next, and a bicycle a week later, they might associate knowing how a car works with knowledge about vehicles more broadly. This inferential strategy could help explain some of the other findings in the current studies: rarely would mechanistic knowledge about one domain be associated with another in practice. After all, someone who can fix a clock is no more likely to demonstrate knowledge about something unrelated, such as tulips, than someone who cannot fix a clock. Additionally, it is easier to observe mechanistic knowledge about artifacts than animals. Artifact knowledge is generally revealed by fixing or modifying an object, actions that children are likely to observe before elementary school. In contrast, mechanistic knowledge about animals is often demonstrated in medical or advanced pedagogical contexts, which children usually do not encounter at length until late in school, if at all.

Another factor guiding children's generalization inferences may be more fundamental intuitions about what kinds of processes or traits category members share. Preschoolers understand that the insides of objects are vital to both their functioning and identity (Gelman & Wellman, 1991). While our stimuli were equated for mentioning internal components, only mechanistic knowledge focuses on their role in causal processes. By Kindergarten, children group processes into domains by considering underlying causal mechanisms instead of relying on surface features (Erickson et al., 2010). In turn, children's sense of what causal processes categories share may guide their epistemic inferences, leading them to generalize knowledge about those mechanisms. In practice, children's epistemic experience and their intuitions about shared properties likely interact to construct our generalization intuitions.

6.2. Limitations and future directions

One potential limitation of our study concerns equating the "strength" of mechanistic and non-mechanistic stimuli. Perhaps participants interpreted the mechanistic twin as acquiring strictly more information, or the mechanistic knowledge as intrinsically more complex. To address this concern, we designed the stimuli to be as matched as possible along multiple dimensions and by using a broad range of factual non-mechanistic knowledge. In addition, all non-mechanistic knowledge examples concerned unobservable traits such as history (e.g. where the first car engines were built) and constitution (e.g. what kinds of glass smartphone screens are made out of). If a particular feature was mentioned in a mechanistic example, it was also mentioned in the corresponding non-mechanistic example to minimize a bias for knowledge about internal parts. The vignettes were also explicitly labeled as example pieces of knowledge, meant to broadly indicate the kind of knowledge each twin possessed rather than specify it exactly. Most importantly, no group of participants uniformly preferred the mechanistic twin, endorsing the mechanistic twin for the superordinate level category significantly more than for the subordinate (older children and adults) or unrelated (all age groups) level category in Study 1a. In Study 1b, adults demonstrated the same pattern. However, older children displayed the reverse pattern, while younger children generally held no strong preference overall. Thus, the relation between knowledge type and generalization was limited to particular categorical levels, arguing against a uniform bias for one stimulus set.

A second potential limitation relates to our use of forced choice judgments. This approach is unable to distinguish whether participants judged the mechanistic twin as knowing more about the superordinate level category because they thought mechanistic knowledge generalizes substantially more, or because they thought non-mechanistic knowledge simply does not generalize at all. Future studies using rating scales could disambiguate these possibilities by providing independent measures for each knowledge type. However, given that the goal of the current studies was to ascertain whether children and adults generalize mechanistic and non-mechanistic knowledge differently, this limitation serves more as a direction for future research.

Another direction for future research might utilize more ecological stimuli. Here, we created carefully matched stimuli that either had a clear focus on mechanism or were completely devoid of mechanism. However, mechanistic and non-mechanistic information are rarely completely separated, either theoretically or cognitively, in the real world. Mechanistic explanations necessarily reference components of a system that have features unrelated to the mechanism, and almost all components of a system will have features that are relevant to various potential mechanistic explanations. Fortunately, contemporary philosophers of science have offered broadly converging accounts of the key features of mechanistic explanations (Bechtel, 2011; Craver & Darden, 2013), and we constructed our stimuli with these in mind.[1]

Future research might also investigate the scope of our generalization intuitions by testing multiple superordinate level categories for a single stimulus. This approach could reveal 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, properties, or epistemic history among 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 aim of the current studies 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 regularity of responses by adulthood, 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 as well as how all of us learn to use this information to better rely on the knowledge of others.

[1] The key features of mechanistic explanation include: 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".

6.3. Conclusion

The current studies were inspired by a growing body of work in cognitive development and the philosophy of science that emphasizes the role mechanism plays in the way all of us, from children to scientists, investigate and evaluate the world. Here, we focused on how mechanism influences our epistemic intuitions by examining whether children and adults selectively generalize mechanistic knowledge. Indeed, even young children recognize that mechanistic knowledge about a basic level artifact category implies greater knowledge about its superordinate level category, compared to factual non-mechanistic knowledge about the same basic level category. In contrast, only adults reliably generalized mechanistic knowledge about biological categories, suggesting a delayed developmental trajectory for that domain. In sum, our studies provide an account of how mechanism influences children and adults' epistemic inferences: mechanistic knowledge signifies rich knowledge about related kinds, an intuition that is demonstrable by the early elementary school years. Over development, these intuitions play a progressively larger role, profoundly shaping our epistemic landscapes into adulthood.

Supplementary materials

All data and materials can be accessed on the Open Science Framework at https://osf.io/fz3vj/.

CRediT authorship contribution statement

Aaron Chuey: Conceptualization, Methodology, Investigation, Formal analysis, Writing - original draft. **Kristi Lockhart:** Conceptualization, Methodology, Writing - review & editing. **Mark Sheskin:** Methodology, Writing - review & editing. **Frank Keil:** Supervision, Conceptualization, Methodology, Writing - review & editing.

Acknowledgments

This research was supported by NSF grant DRL 1561143 awarded to Frank C. Keil.

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.

Buchanan, D. W., & Sobel, D. M. (2011). Mechanism-based causal reasoning in young children. *Child Development, 82*(6), 2053–2066.

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. (2009). Preschoolers' search for explanatory information within adult–child conversation. *Child development, 80*(6), 1592–1611.

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., & Wellman, H. M. (1991). Insides and essences: Early understandings of the non-obvious. *Cognition, 38*(3), 213–244.

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. (2018). Preferences for explanation generality develop early in biology but not physics. *Child Development, 89*(4), 1110–1119.
Kaiser, M. I., & Krickel, B. (2017). The metaphysics of constitutive mechanistic phenomena. *The British Journal for the Philosophy of Science, 68*(3), 745–779.
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., Vredenburgh, 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. (2019). The privileged status of knowing mechanistic information: An early epistemic bias. *Child Development., 90*(5), 1772–1788.
López, A., Gelman, S. A., Gutheil, G., & Smith, E. E. (1992). The development of category-based induction. *Child Development, 63*(5), 1070–1090.
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.
NGSS Lead States (2013). Next generation science standards: For states, by states. Retrieved from https://www.nextgenscience.org.
Sheskin, M., & Keil, F. C. (2018). *TheChildLab.com: A video chat platform for developmental research.* psyarxiv https://psyarxiv.com/rn7w5.
Stavy, R., & Wax, N. (1989). Children's conceptions of plants as living things. *Human Development, 32*(2), 88–94.
Taborda-Osorio, H., & Cheries, E. W. (2017). Developmental origins of biological explanations: The case of infants' internal property bias. *Psychonomic Bulletin & Review, 24*(5), 1527–1537.
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. (2020). *Why teach how things work? Hidden benefits of exposure to mechanistic explanation.* (under review).
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.