

**Mechanistic Reasoning in 3- to 6- year-old Children: Investigating the Role of
Failure and Explanation**

A thesis submitted by

Gauri Harindranath

in partial fulfillment of the requirements for the degree of

Master of Science

in

Psychology

Tufts University

May 2024

© 2024, Gauri Harindranath

Adviser: Paul Muentener

Abstract

We investigated whether contexts of failure and providing explanations improved children's exploration, explanation, and generalization of mechanistic knowledge. In Study 1, we showed 3- to 6-year-old children how a toy works and manipulated whether a hidden switch was ON (Success) or OFF (Failure). Children in the Failure condition were more likely to discover the switch and talk about causal mechanisms. Older children were more likely to generalize mechanistic learning to a novel context. In Study 2, we investigated generalization in 3- to 4-year-olds by turning the switch OFF and providing participants with no explanation about the toy, an explanation before exploring the toy, or an explanation after exploring the toy. There were no significant differences in generalization across conditions, but children who explored longer were more likely to generalize information when they received an explanation. The findings discuss the role of failure and explanations in children's causal learning.

Keywords: causal reasoning, mechanistic reasoning, failure, exploration, explanation, generalization

Table of Contents

Introduction	5
Causal reasoning and ambiguity	5
Mechanism reasoning across childhood	7
Causal failure and mechanism reasoning	10
Study 1	13
Method	13
Participants	13
Materials	14
Procedure	15
Coding	17
Deviations from preregistered procedure and analysis	18
Results	19
Preliminary analysis	19
Exploration Phase	20
Explanation Phase	22
Generalization Phase	23
Discussion	25
Study 2	26
The role of explanations	26
Timing of explanations	28
Method	31

Participants.....	31
Materials.....	32
Procedure.....	32
Coding.....	34
Preregistered procedure and analysis.....	34
Results.....	34
Exploration Phase.....	34
Generalization Phase.....	36
Exploratory Analyses.....	38
Discussion.....	40
General Discussion.....	42
Role of failure in mechanistic reasoning.....	43
Relation between exploration, explanation, and generalization.....	44
Conclusion.....	47
References.....	48

Introduction

Children have demonstrated a capacity for reasoning about the relation between causes and effects, along with thinking about an event's underlying structures (Bonawitz et al., 2012, Schulz & Sommerville, 2006, Muentener & Bonawitz, 2017, Gweon & Schulz, 2008, Buchsbaum et al., 2012, Bonawitz et al., 2010). Starting at an early age, children engage in causal learning – often by generating hypotheses, testing them out by collecting evidence, and drawing conclusions based on this evidence – processes involved in the scientific process (Muentener & Bonawitz, 2017; Schulz, 2012). One example of a causal relation is flipping a switch (cause) to turn on a light (effect). Mechanisms, such as electricity, enable the cause of flipping the switch to elicit the effect of turning on the light. Mechanistic reasoning – how and why systems work the way they do – involves understanding the visible (e.g., batteries) and invisible entities (e.g., electricity, force) that link causes to effects. While we have evidence that children are capable of reasoning about mechanisms that support the relation between a cause and its effect (Buchanan and Sobel, 2011), we know less about the contexts that spontaneously promote mechanistic reasoning. Thus, investigating how mechanistic reasoning changes over development, and the contexts that support such reasoning, can have implications for an understanding of children's early scientific reasoning abilities. In this project, we consider whether causal failure is one of the contexts that support mechanism reasoning, and if providing mechanistic explanations helps children generalize what they learned from the context of causal failure to novel situations.

Causal reasoning and ambiguity

Children can start reasoning about causal relations by at least the preschool years and having knowledge about these relations allows us to make predictions about future events and

generate interventions to bring about specific events (Gopnik et al., 2004). Children's actions are motivated by causal determinism – a belief that all effects have causes (Muentener & Schulz, 2014; Schulz et al., 2008; Schulz & Sommerville, 2006). If actions reliably predict outcomes, they infer that the actions are causing the outcome (Gopnik et al., 2001; Gopnik & Schulz, 2004; Gopnik & Sobel, 2000; Kushnir & Gopnik, 2005; Nazzi & Gopnik, 2000; Schulz & Gopnik, 2004).

Even when viewing unpredictable causal relations, children have a strong sense of causal determinism. Schulz and Sommerville (2006) demonstrated a causal relation to 3- to 6-year-old children that occurred either deterministically or probabilistically. After the demonstration, experimenters would reveal an object in their hand without stating whether or not it impacted the causal relation. When children were asked to generate an effect based on their knowledge of the causal relation, those in the probabilistic condition were more likely to interact with the revealed object. This indicates that uncertainty in causality influences children's understanding of causal relations and they can infer the role of hidden causes.

When causal relations are unclear, children are motivated to understand them better. Schulz and Bonawitz (2007) showed 4- to 6-year-old children a toy with either an obvious causal relation (where one cause led to an effect) or an ambiguous causal relation (where there were two potential causes leading to an effect). When children were given the opportunity to choose between exploring the demonstrated toy or a novel toy, those who viewed the ambiguous causal relation were more likely to choose the demonstrated toy. Similarly, Muentener and Schulz (2014) show that when viewing an unfamiliar machine light up, toddlers are more likely to explore the machine when it appears to light up on its own compared to when they already know how it works.

This work collectively indicates that children are capable of exploring causal relations and enhanced exploration due to causal ambiguity may lead to higher levels of mechanistic reasoning. However, whether they are learning about the mechanisms through this exploration is less clear. Russ et al. (2008) argue that a bulk of causal research does not actively consider the underlying mechanistic structure that supports the cause-and-effect relation and there has not been sufficient demonstration of spontaneous mechanistic reasoning in children. In the current study, we address these concerns by assessing the extent to which children's exploration, explanation, and generalization include mechanistic reasoning in particular.

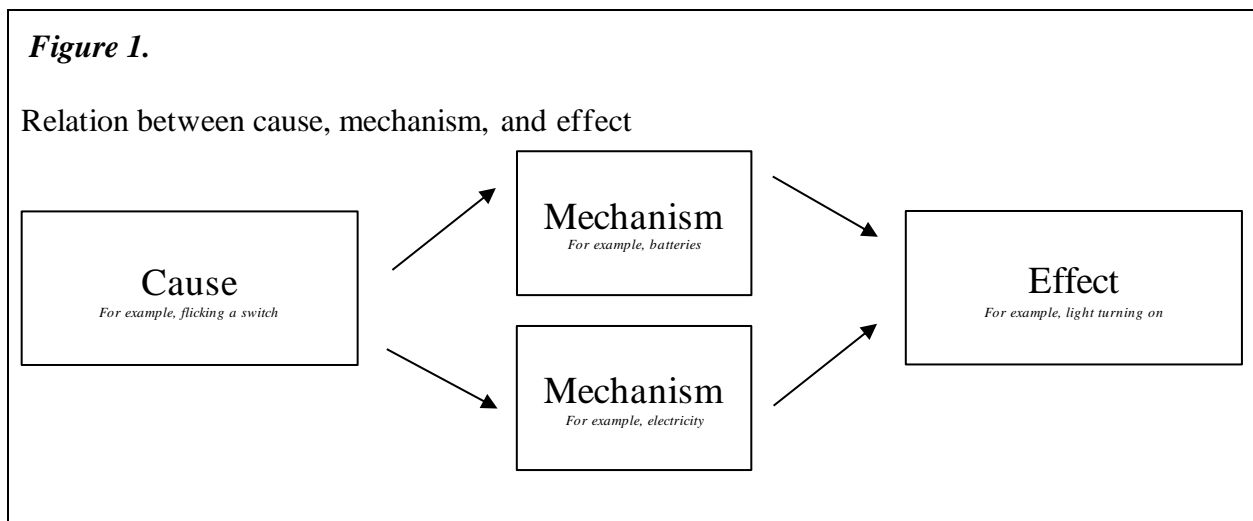
Mechanism reasoning across childhood

We define mechanistic reasoning as any behavior, exploratory or explanatory, that goes beyond the demonstrated features of a causal event towards content in the space between actions and outcomes. While there has been more experimental work related to mechanisms in causal relations in the past decade (Buchanan & Sobel, 2011, Stephan & Waldmann, 2022, Ahl & Keil, 2017, Erb & Sobel, 2013), the contexts that promote children's spontaneous mechanism reasoning is less clear.

Children understand the role of mechanisms in a causal relation and can make predictions about mechanistic interventions (e.g., what happens when you remove a battery from a remote). Buchanan and Sobel (2011) showed 3- and 4-year-olds three circular lights – two small and one large – placed on a board. Initially, they connected one small light to a large one using a wire and demonstrated a causal relation, where pushing the small light turns on the large light. Subsequently, they switched the wire to connect the other small light to the large light. Four-year-olds (compared to 3-year-olds) were more likely to infer the role of the wire and how an

intervention of switching the wire placement impacted the relation between the small and big lights.

Buchanan and Sobel (2011) showed that understanding of mechanisms improves with age and exposure. They hypothesized that the limited performance of 3-year-olds in the aforementioned study may not be due to an inability to reason about mechanisms but rather a lack of familiarity with wires. They replicated their previous study with batteries, a more familiar mechanism to 3-year-olds, instead of wires and found an improvement in 3-year-olds' performance. These findings indicate that, based on exposure, children as young as 3-years-old are able to reason about mechanisms. Figure 1 demonstrates the relation between cause, effect, and mechanisms.



While we have evidence for children's ability to reason about mechanisms in experimental contexts, research investigating children's spontaneous mechanism reasoning has yielded mixed results. Some studies show that when preschool-aged children are prompted to explain how causal systems work, they are more likely to learn about mechanisms through their subsequent exploration. For instance, children who provide explanations related to causal

functions (i.e., why an object did not work) display behaviors more consistent with searching for hidden mechanisms (Legare, 2012). However, it is unclear whether such exploration actually leads to enhanced mechanistic learning.

Other studies show that when children are provided with mechanistic explanations about a causal system, children may be better able to learn about mechanisms (Kurkul et al., 2021). Four- and 5-year-olds interacted with a circuit board that connected a switch (cause) to a toy fan (effect) with mechanisms such as batteries and wires in the middle. The experimenter would ask questions about the board (such as “*Why is there a battery?*”) and children either heard non-mechanistic responses (such as “*To make it work?*”) or mechanistic responses (such as “*The battery is there to give the toy power. When the batteries are connected to these buttons, then the toy will work because it has power?*”) in response to the questions. The findings revealed that children provided with mechanistic explanations were more likely to reconstruct the circuit board. This suggests that children’s mechanism reasoning can benefit from hearing mechanistic explanations but there is room to explore how children independently may be prompted to explore mechanisms.

However, other studies show that similar cues may not improve children’s mechanistic reasoning. For instance, Attisano et al. (2021) investigated whether prompting children to think about mechanisms would influence their mechanistic reasoning. Four- to 8-year-olds were either given a control prompt about the history of a machine (“*It has worked the way it does for a long time?*”) or a mechanistic prompt (“*The parts inside it make it work the way it does?*”). Parent-child dyads were then allowed to freely explore and talk about the exhibit, after which, children were asked questions about the machine. The researchers found that children’s discussion of the machine and its components were correlated with their parents’ talk about the

same but their mechanistic talk was not influenced by parents' mechanistic talk. Additionally, they found no effect of the prompts on the dyads' conversations or children's responses to questions about the machine. Overall, directing parent-child dyads to consider *how* a causal system works improved children's knowledge of the system but did not lead to an increase in their mechanism-related talk and knowledge. Similar research showed that conversation prompts encouraging parents to *explore* or *explain* a museum exhibit led to heightened levels of exploration and explanation, respectively, in dyads (Willard et al., 2021). However, this did not extend to children's mechanism learning. Our project seeks to examine the role of causal failure in mechanistic reasoning and how it might contribute to explaining the mixed results discussed above.

Causal failure and mechanism reasoning

Causal failure motivates different patterns of exploration in children, starting in infancy. For instance, infants as young as 16-month-olds make judgments about causal failure and adjust their actions accordingly (Gweon and Schulz, 2011). Infants were shown a causal relation, where pushing a button on a toy would play music. Infants were then either given the same toy (within-object condition) or a novel but identical one in a different color (between-object condition). Neither toy would work when infants pushed the button but, depending on the condition, they demonstrated different patterns of behavior. In the within-object condition, infants were more likely to hand the toy to their parent, indicating their understanding that their actions may be the potential source of failure. Conversely, in the between-object condition, infants were more likely to reach for the original toy, indicating their understanding that the novel toy may be the potential source of failure. These findings suggest that infants are able to hypothesize about failure and implement appropriate interventions to resolve it.

When causal systems are working, children could display curiosity about their mechanisms (e.g., learning that batteries provide power to remote controls). However, knowledge of mechanisms becomes critical when causal systems fail to work (e.g., one needs to replace batteries when they fail). Legare (2012) showed that when 2- to 6-year-old children who see inconsistent causal relations (where cause does not lead to the anticipated effect) compared to consistent causal relations are more likely to provide “*causal function*” explanations – explanations that focus on the internal parts of the causal systems. Furthermore, children who viewed inconsistent causal relations provided more *causal function*-based explanations than other types of explanations (such as *causal action* explanations, that focus on the functioning of the system). Children who provided *causal function* explanations spent more time engaging with the causal system than those who provided other types of explanations. This indicates that children viewing inconsistent causal relations are more likely to talk about the internal properties of the causal system (such as mechanisms), leading to a higher level of exploration of the system.

Similar to Legare (2012), several studies in the area of causal or mechanism reasoning incorporate elements of failure. For instance, they may create situations in which children are presented with evidence – shown by an experimenter or generated on their own – that an action inconsistently causes an outcome to occur (e.g., probabilistic evidence) and then assess their exploratory behaviors (e.g., Gopnik et al., 2001; Schulz & Bonawitz, 2007). The results of these studies suggest that inconsistency or ambiguity, a form of causal failure, motivates a different pattern of exploratory behavior.

Additionally, upon further inspection, Kurkul and colleagues (2021) included failure-oriented language in their mechanistic explanations (e.g., “The toy would not work because the snaps need to be touching to move the power...”), which may have facilitated children’s

mechanistic reasoning beyond the success-oriented language also included in the study. In contrast, the prompt and exhibit used by Attisano and colleagues (2021) informed children that “the parts inside of it [a machine] make it work the way it does” and asked them to “learn what you can about the machine”, information that at least implicitly focuses on the successful use of the machine. Finally, although Willard and colleagues (2019) found that explanation prompts did not lead to improved mechanistic reasoning, they found that children who spontaneously engaged in troubleshooting behavior (i.e., experienced failure) had better memory and understanding of causal mechanisms.

Overall, we have ample evidence that children explore more following contexts that involve failure, but we have less evidence that this exploration leads to mechanistic learning and whether any learning extends to new contexts. Thus, in Study 1, we investigated whether causal failure, contrasted with causal success, is a context that promotes a child’s discovery and learning about causal mechanisms. Children were given causal instructions about a toy to use for a goal-directed task. They were not told about a partly hidden switch (i.e., a mechanism) that needed to be turned ON for the toy to work. Using a between-subjects design, we manipulated whether the switch was ON or OFF (i.e., Success vs. Failure) before handing it to children. We assessed children’s exploration of the toy, their explanations for how it worked, and any learning generalized to a novel but similar toy. We predicted that if causal failure leads to improved mechanistic reasoning, children who experience failure would, 1) spend more time exploring the toy, 2) be more likely to find the hidden mechanism, reference the mechanism in their causal explanations, and 3) generalize their learning by looking for the mechanism in a novel but similar toy.

Study 1

Method

Participants

Seventy-one preschoolers ($M = 5.13$ years; range: 3.14 - 6.99 years) were recruited at an urban science museum. To ensure a sufficient number of children across the entire age range, we recruited children into separate two-year age bins (3- and 4-year-old children, $N = 34$, $M = 4.1$ years; 5- and 6-year-old children, $N = 37$, $M = 6.0$ years). Following standard museum protocols, caregivers and legal guardians were not required to provide demographic information.

Caregivers of 46 participants opted to report the child's demographic information. Out of those 46 participants, 27 were male and 19 were female. Additionally, 32 participants identified as White, 3 as Hispanic/Latine, 2 as Black, 6 as Asian, and 4 as multiracial individuals.

Additionally, 44 caregivers shared their highest level of education – 3 adults finished some college, 13 adults finished college, and 28 adults had an advanced degree. The demographics of this subsample of participants matches those expected from the typical museum visitor population, which is typically composed of White children from middle- to high-SES families. Although research has shown that children across distinct cultures and socioeconomic statuses exhibit similar causal reasoning abilities (Wente et al., 2019), we acknowledge that the current sample may nonetheless warrant relative caution when attempting to generalize any study results to the broader population.

An additional 17 participants were recruited but not included in the final sample for the following reasons: procedural error ($N = 2$), technology issue (e.g. video recorder not working) ($N = 1$), and caregiver interference ($N = 14$).

Participants were randomly assigned to one of two conditions: 34 children in the Success condition and 37 children in the Failure Condition. Post-hoc power analyses, conducted using

G*Power software, suggest that this was a sufficient sample size to detect a significant difference between the Failure and Success conditions for a logistic regression analysis on the binary outcome variables discussed below. We conducted these power analyses on a range of potential baseline levels in the Success condition, at least a 20% increased likelihood of responding in the Failure condition, and power of at least .8.


Materials

Toys from MindWare Start-Up Circuits were modified and used as stimuli in this study (Figure 2). For the *Exploration* and *Explanation* phases described below, we used a fan toy. The fan included two parts (upper fan component and lower switch component) and worked when the two parts of the toy were connected and the switch was in the ON position. We modified the toy by adding a cardboard box to cover the sides of the lower component so that the switch was not visible from a side angle (Figure 2A). However, the bottom of the lower component was not covered so the switch was visible from that angle and could easily be discovered upon manipulating the toy (Figure 2B). We also added a red line on the upper and lower components to help children correctly align the pieces. Thus, during exploration, the child could discover that the fan worked when the two parts were put together so that the red lines touched and the switch was in the ON position.

For the *Generalization* phase, we used the light toy from MindWare Start-Up Circuits. Similar to the fan toy, the light worked when the two pieces of the toy were connected and the switch was in the ON position. We modified the toy by adding cardboard to the lower component so that the switch was not visible from the side (Figure 2C) but could be seen from the bottom (Figure 2D).

Figure 2

Stimuli used in Exploration, Explanation, and Generalization Phases



Note: Top section: Fan toy used during *Exploration* and *Explanation* Phases. A) View of the two parts of the toy consistent with how participants received the toy at the start of the *Exploration* phase. B) View showing the switch that was partly hidden from participant's initial view. Note that participants could not see the switch when the toy was presented to them, but it could be found easily upon manipulating the toy. Bottom section: Light toy used during the *Generalization* Phase. C) View of the two parts of the toy consistent with how participants received the toy at the start of the *Generalization* phase. D) View showing the switch that was partly hidden from participant's initial view. Similar to the fan toy, participants could not see the switch when the toy was presented to them, but it could be found easily upon manipulating the toy.

Procedure

Each participant participated in four phases of the study: an *Introduction* phase, an *Exploration* phase, an *Explanation* phase, and a *Generalization* phase. All procedures were approved by the Tufts Institutional Review Board.

Introduction Phase

All participants were first shown how the fan toy worked and a game that they could play with the fan toy. The experimenter told the participant, “*I just got this toy and I read that it works when you put the two parts together so that the red lines touch*”. This was accompanied by the experimenter simultaneously putting the toy together correctly and showing the participant that the toy worked. Participants were told they could use the toy to play a game in which they blew slips of paper into a box hanging from the edge of the table.

Exploration Phase

Participants were told that the toy used for the demonstration would stay with the experimenter but an identical version of the toy was placed on the table in front of them. Between conditions, we manipulated whether the toy would work for the participant by surreptitiously turning the switch at the bottom of the lower component ON (Success Condition) or OFF (Failure Condition) before giving the toy to the participant. Thus, in the Success condition, when participants put the toy together with the red lines aligning, the fan would turn on. In contrast, in the Failure condition, given the switch was turned OFF, the fan would not turn on when the participants put the toy together with the red lines aligning.

Participants were handed the two components of the toy separately and allowed to freely explore the toy and play the game for up to 5 minutes or when they said they were done. If the participant did not interact with the toy for 10 seconds, the experimenter asked them whether they wanted to keep going or were done. The *Exploration* phase ended when the participant said they were done exploring or they reached the 5-minute maximum allowed exploration time.

Explanation Phase

After the *Exploration* Phase, the toy was kept on the table in front of the participant. To elicit participants’ spontaneous causal understanding of the toy and mechanism-related language,

participants in both conditions were first asked, “*What can you tell me about the toy?*” To elicit more directly their causal understanding and mechanism-related language, participants were next asked, “*How does it work?*”

Generalization Phase

Finally, the experimenter told the participant that they had another toy (i.e., the light toy) and that they did not know how it worked. The experimenter asked the participant “*Can you help me figure out how it works?*” while giving them the novel but similar light toy. The experimenter then asked them “*How does it work?*” The first question aimed to get participants to engage with the toy and the second question aimed to elicit explanations. Critically, the switch in this toy was in the ON position for participants in both conditions. Thus, upon connecting both components, the toy would work for participants in both the Failure and Success conditions. Providing participants with a working toy was necessary to ensure that we were assessing participants’ generalization from the *Exploration* and *Explanation* phases, rather than newly discovered causal mechanisms following failure if they were to have been provided with a non-working toy.

Coding

Each participant’s exploratory behaviors, verbal explanations, and nonverbal gestures were transcribed and coded by two research assistants, unaware of the participant's condition. Agreement on the continuous exploratory play measure was high, $r > .95$. Disagreements during the coding of the binary exploration and explanation measures were infrequent and subsequently decided by a third researcher.

Exploration Phase

We transcribed participants’ exploratory behaviors across the entire *Exploration* phase. Our primary dependent variables of interest included the proportion of time participants directly explored the toy out of their total time exploring the toy and playing the goal-directed game. We

also coded whether participants physically interacted with the switch by turning it ON or OFF (coded in a binary manner: 0 = No, 1 = Yes). Furthermore, after data collection began, we noted that some participants had difficulty aligning the red lines correctly, and therefore decided to code this behavior as well.

Explanation Phase

We coded participants' responses for whether they verbally or non-verbally (e.g., pointing) referenced the following parts of the toy: connecting the two parts, the switch, and general mechanisms (batteries, electricity, and wires) (coded in a binary manner: 1 = Yes, 0 = No).

Generalization Phase

We coded whether participants referenced the switch (verbally or non-verbally) (coded in a binary manner: 1 = Yes, 0 = No).

Deviations from preregistered procedure and analysis

All coding measures and predictions are registered on Open Science Framework (OSF) (https://osf.io/7p3tc/?view_only=e0eda574354e483183f3f833180d7cdf). We made three deviations from this preregistration. First, given the primary focus of the study was on whether children engaged in mechanistic reasoning during contexts of failure, rather than the extent to which they did so, we coded only for the presence of mechanistic utterances as a binary measure rather than the frequency of mechanistic utterances, as was initially noted in the preregistration. Second, although we intended to analyze children's reference to other mechanisms (e.g., batteries, wires, etc), such references were very infrequent – used by only 8 out of 71 participants – and not included in our final analyses. Finally, instead of using age and condition to predict whether or not children referred to the switch in the *Generalization* phase, we used age and whether children found the switch in the *Exploration* phase as predictors in our regression

model. As discussed below, we chose not to use condition as a predictor given few children found the switch in the Success condition (1/18 participants).

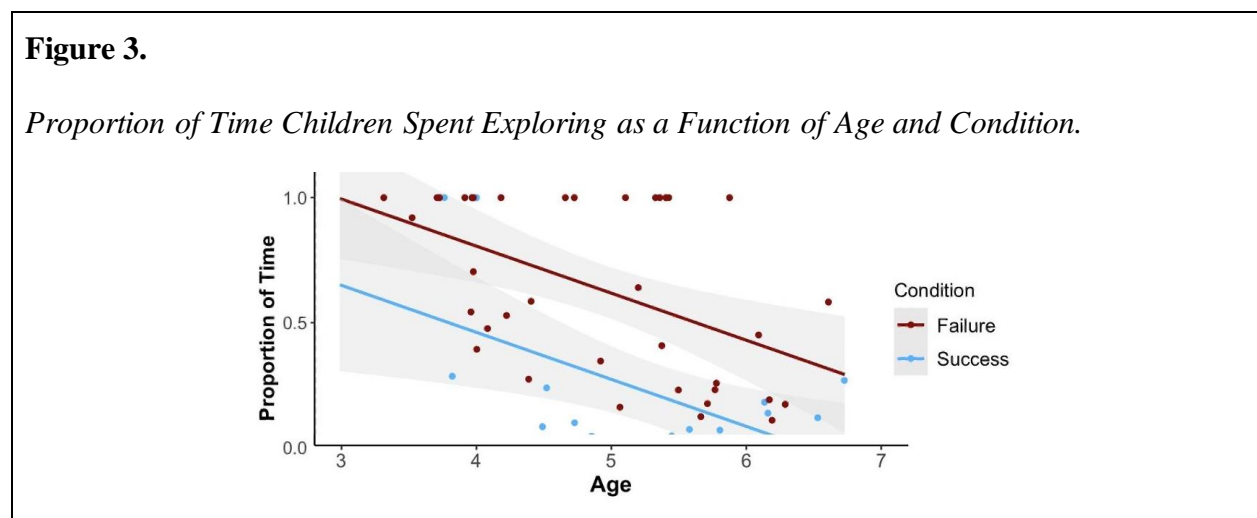
Results

Preliminary analysis

All participants, in both conditions, received an instruction to put the two parts of the toy together so that the red lines on both parts aligned. These instructions were intended to set up the primary condition difference between success and failure: Children in the Success condition would succeed in making the toy work (i.e., it would turn on when the parts were connected) while children in the Failure condition would fail to make the toy work (i.e., the toy would not turn on when the parts were connected). However, we found that 29 participants (13 in the Failure condition, 16 in the Success condition) did not align the two parts of the toy according to our instructions during their first interaction with the toy. As a result, and counter to our intended manipulation, a subset of participants in the Success condition experienced failure, given that the two parts had to be aligned correctly for the toy to work. To ensure that our analysis of participants' mechanistic behaviors below focused on the intended contrast between participants who experienced failure compared to those participants who experienced success, we omitted the 16 participants in the Success condition from the subsequent analyses. Thus, all data from participants in the Success condition reported below experienced success on their first interaction with the toy. Given that participants in the Failure condition could not initially make the toy work, regardless of how they aligned the parts, we did not omit their data from the subsequent analyses. Separate analyses including all data from all participants in both conditions produce the same pattern of results as those reported below.

Exploration Phase

First, we assessed whether participants were more likely to explore the toy (relative to the goal-directed game) when they experienced failure. We conducted a linear regression to determine if condition predicted the proportion of time participants spent exploring the toy. Condition (Failure vs. Success) was entered as a fixed factor, and Age was entered as a continuous predictor. We found that Age was a significant predictor of the proportion of time participants spent with the toy, $t = -4.36, p < 0.001$ (Figure 3).



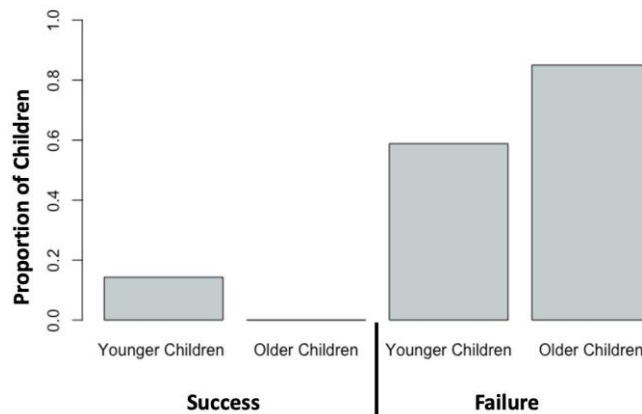
Contrasting younger (3- and 4-year-old) participants with older (5- and 6-year-old participants), showed that younger participants ($M = 0.67, SD = 0.36$) spent a higher proportion of time exploring the toy than older participants ($M = 0.36, SD = 0.35$). Condition was also a significant predictor, $t = -4.08, p < .001$; participants spent more time exploring the toy when they experienced failure ($M = 0.64, SD = 0.35$) than when they experienced success ($M = 0.21, SD = 0.30$). A subsequent model including the interaction between Condition and Age was not significant, $t = 0.001, p = 0.999$.

Next, we assessed whether participants were more likely to discover the switch when they experienced failure. We conducted a logistic regression on participants' likelihood to

interact with the switch (0 = No, 1 = Yes) with condition (Failure vs. Success) entered as a fixed factor, and age was entered as a continuous predictor. Condition was a significant predictor with participants in the Failure condition being more likely to find the switch (27/37 participants, 72.97%) than participants in the Success condition (1/18 participants, 5.56%), $z = -3.55, p < 0.001$ (Figure 4). Age was not a significant predictor, $z = 0.98, p = 0.327$. A subsequent model including the interaction between Condition and Age was not significant, $z = -1.23, p = 0.218$.

Figure 4.

Proportion of Children Finding the Switch during Exploration as a Function of Age and Condition.



As an additional exploratory measure, we also assessed whether the amount of time participants spent exploring the toy predicted their interaction with the switch. We conducted a logistic regression on participants' likelihood to interact with the switch (0 = No, 1 = Yes) with condition (Failure vs. Success) entered as a fixed factor, and exploration time was entered as a continuous predictor. Condition was a significant predictor with participants in the Failure condition being more likely to find the switch, $z = -3.67, p < 0.001$. Exploration time was not a

significant predictor, $z = -1.20$, $p = 0.228$. A subsequent model including the interaction between Condition and Exploration Time was also not significant, $z = 0.007$, $p = 0.995$.

Explanation Phase

In the *Explanation* phase, we asked children “*What can you tell me about the toy*” and “*How does it work?*”. Across both questions, 15 participants (6 in Failure and 9 in Success) referenced how to put the toy together (referencing our initial instructions), 15 participants (14 in Failure and 1 in Success) referenced the switch, and 7 participants (6 in Failure and 1 in Success) referred to putting the toy together and the switch. Furthermore, 10 participants (3 in Failure and 7 in Success) gave us explanations such as “*the fan spins*” or “*it blows stuff*” and 8 participants (all in Failure) did not respond or told us they did not know. Along with one of the aforementioned explanations, 4 participants (1 in Failure and 3 in Success) referenced general mechanisms (such as batteries or electricity).

We focused on whether participants were more likely to spontaneously reference the hidden switch when they experienced failure. We conducted a logistic regression on participants’ likelihood to spontaneously reference the switch (binary, 0 = No, 1 = Yes) in their response to the first explanation question (“*What can you tell me about the toy?*”). Condition (Failure vs. Success) was entered as a fixed factor and Age was entered as a continuous predictor. Neither Condition, $z = -0.75$, $p = 0.453$, nor Age, $z = 0.71$, $p = 0.475$, were significant predictors. A model with interactions was also not significant, $z = 0.025$, $p = 0.98$.

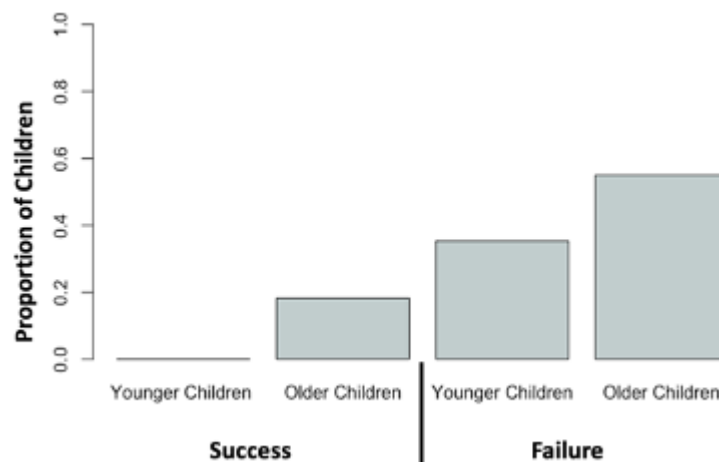
We then conducted a logistic regression to determine if the condition (Failure vs. Success) predicted whether or not participants ever referenced the hidden switch during the *Explanation* phase. For this analysis, we conducted a logistic regression on participants’ likelihood to mention the hidden switch (binary, 0 = No, 1 = Yes) in their response to either question (“*What can you tell me about the toy?*” and “*How does it work?*”). Condition (Failure

vs. Success) was entered as a fixed factor and Age was entered as a continuous predictor.

Condition was a significant predictor with participants in the Failure condition being more likely to mention the switch (17/37 participants, 45.95%) than participants in the Success condition (2/18 participants, 11.12%), $z = -2.56$, $p = 0.011$ (Figure 5). Age was not a significant predictor of participants mentioning the switch, $z = 1.49$, $p = 0.137$. A model with interactions was not significant, $z = 0.43$, $p = 0.67$.

Figure 5.

Proportion of Children Referencing the Switch during Explanation as a Function of Age and Condition.



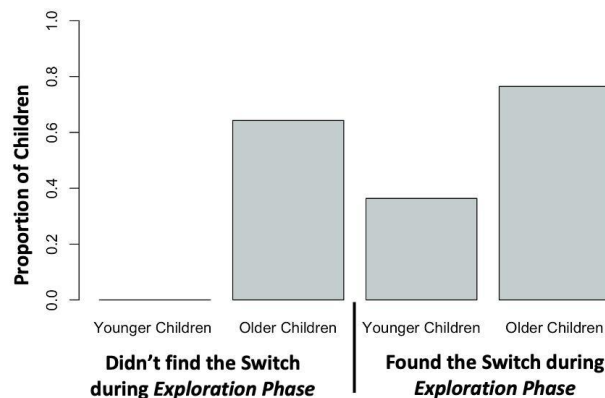
Generalization Phase

Finally, we assessed whether children generalized any mechanistic reasoning. A logistic regression was conducted on children's likelihood to reference the switch (binary, 0 = No, 1 = Yes) in their responses about the toy; finding the switch (0 = No vs. 1 = Yes) during the *Exploration* Phase was entered as a fixed factor, with Age was entered as a continuous predictor; given that children rarely discovered the switch in the Success condition (only 1 of 18 participants), we did not include Condition as a fixed factor. Analyses revealed that Age was a

significant predictor of children referencing the switch during the *Generalization* phase, $z = 3.11$, $p = 0.002$ (Figure 6); older children (22/31 participants, 70.97%) were more likely to reference the switch than younger participants (5/24 participants, 20.83%). Finding the switch was also a significant predictor, $z = 1.98$, $p = 0.048$; children who found the switch during the *Exploration* phase were more likely to find the switch during the *Generalization* phase (17/28 children, 60.71%) than children who did not find the switch in the *Exploration* phase (10/27, 37.03%). A subsequent model with interactions was not significant, $z = -0.81$, $p = 0.419$.

Figure 6.

Proportion of Children Referencing the Switch during the Generalization Phase as a Function of Age and Whether they Found the Switch during exploration.



Breaking down the children who found the switch during the *Exploration* phase by age, 4 out of the 11 (36.37%) 3- and 4-year-olds who found the switch in the fan toy went on to find the switch in the light toy. In contrast, 13 out of the 17 (76.47%) of 5- and 6-year-old children who found the switch in the fan toy went on to find the switch in the light toy. These analyses show that children were more likely to reference the switch during this phase if they had previously discovered the switch – consistent with the aim that this phase assessed children's generalization

of their learning – but that children’s overall ability to generalize their mechanistic reasoning increased with age.

Discussion

The current study investigated whether contexts of failure support children’s mechanistic reasoning. We gave 3- to 6-year-old children a toy along with instructions on how to use it to play a goal-directed game. Between conditions we manipulated whether a partly hidden switch was turned off, leading children to either succeed or fail to make the toy work. We then measured whether children discovered the mechanism, whether they included it in their explanations about the toy, as well as whether they generalized any exploration and explanation behavior to a new, but similar, functioning toy.

This work was motivated by research that illustrates a child’s ability to reason about causal mechanisms. We aimed to highlight causal failure as a condition in which children are capable of higher levels of independent exploration and reasoning related to mechanisms. As predicted, children in the Failure condition (where the switch was OFF) were more likely to interact with the switch than children in the Success condition (where the switch was ON). Moreover, children in the Failure condition were also more likely to reference the switch when talking about the toy. Age predicted the proportion of time children engaged with the toy – older children proportionally spent less time with the toy than younger children, yet they were more likely to interact with the switch.

In the case of a novel but similar toy, we did not find an effect of condition in whether or not children referenced the switch. However, we did see an effect of age, with older children being more likely to reference the switch than younger children. The absence of a condition effect may be attributed to a few factors. For instance, unlike the game children could play with the fan toy, and there was no goal-directed task for children to perform with the light toy.

Furthermore, children were instructed to “*figure out how the toy works*” with the light toy, which may have prompted a different pattern of exploration when compared to the fan toy. Further investigation is required to understand these results, along with the relation between exploration, explanation, and generalization.

Our findings suggest that causal failure may lead to children discovering and referencing the mechanisms relevant to failure. However, older children in the current study were better able to generalize their mechanistic learning to novel contexts. Specifically, among the 24 3- and 4-year-old participants, 11 (1 in Success and 10 in Failure) found the switch in the fan toy. Yet, only 4 participants referenced the switch in the light toy. This is compared to 31 5- and 6-year-olds, 17 (0 in Success and 17 in Failure) of whom found the switch in the fan toy: 13 children referenced the switch in the light toy. Overall, we see that around 36.37% of younger children generalized their findings of the switch to the novel toy compared to 76.47% of older children.

These results indicate that younger children are less likely to generalize learned mechanistic knowledge than older children. To further explore the factors that might support children’s generalizations, the next study investigated the role of mechanistic explanations, which include references to mechanisms and their functions. Furthermore, we considered whether the timing of these explanations (providing them before or after children’s exploration) impacted children’s ability to generalize learned mechanistic knowledge.

Study 2

The role of explanations

Generalization involves the transfer of learned knowledge to a novel context. In this study, we investigated whether children spontaneously generalized the acquired knowledge about the fan switch to the novel but similar light switch. In Study 2, we focused on the *Generalization* phase where children were given a novel but similar toy. The switch (i.e., the

mechanism) was ON and worked for children since we aimed to study whether they spontaneously applied the mechanistic knowledge they acquired from failure. We investigated the impact of mechanistic explanations and their timing on this generalization.

Explanations have been shown to support a child's learning starting at an early age (Legare, 2012, Leech et al., 2020, Mills et al., 2022, Callanan et al., 2017). Children as young as 3- years old start asking "why" and "how" questions in their everyday conversations to discern causal relations (Callanan & Oakes, 1992). Furthermore, Frazier et al. (2009) studied 2- to 5-year-old children's conversations in naturalistic and experimental settings. They found that when asking "why" and "how" questions, children preferred explanatory responses over non-explanatory ones. When children heard non-explanatory questions, they were likely to respond by re-asking the question or coming up with their own explanations. Along with hearing causal explanations, prompting children to explain causal events focuses their attention on causally relevant details (Walker et al., 2014, Walker et al., 2017).

While it is generally believed that explanatory frameworks with mechanisms should be introduced in secondary school, Kelemen (2019) argues that that is too late – children should be introduced to these explanations by early elementary years. By the age of 5-years-old, children demonstrate a preference for mechanistic explanations over non-mechanistic explanations, considering those who provide mechanistic explanations to be more knowledgeable (Lockhart et al., 2019). Prior research shows that mechanistic explanations may help children generalize learned mechanistic knowledge in different contexts. Lombrozo et al. (2018) demonstrated that when 4- and 5-year-olds are given mechanistic explanations about natural phenomena, they can generalize this mechanistic information to novel situations as well. Furthermore, Kurkul et al. (2021) also investigated how mechanistic vs. non-mechanistic explanations influenced children's

generalization. Children were instructed to assemble a circuit using components that were novel but similar to the ones they had received explanations about. They found that children who received mechanistic explanations were more likely to assemble a working circuit than those who received non-mechanistic explanations. To sum, we know children prefer causal explanations to their questions and their ability to generalize their learning increases with mechanistic explanations.

Based on Study 1, we have evidence that mechanistic exploration increased in the context of failure but younger children were less likely to generalize the information that they learned. The discussed work suggests that if children hear mechanistic explanations about a causal system (i.e., the fan toy), they may be more likely to generalize their learning to a novel situation (i.e., the light toy).

Timing of explanations

Children's exploration is influenced by when they are presented with information. For example, Sobel et al. (2022) investigated parent-child dyads interaction with an electric circuit exhibit. They provided one group of dyads with "initiation" messages (describing the affordances of the exhibit components; e.g., "This is a button"), and the other group with "breakthrough" messages (promoting exploration of the exhibit; e.g., "Can you figure out what else to do?"). After some exploration, each group received the other type of messages and continued exploring. Dyads who received initiation messages first were more engaged with the exhibit, and after their exploration, those children completed more circuit-related challenges. This indicates that providing messages, and perhaps explanations, at different times impacts children's exploration and generalization of knowledge.

In this project, we also aimed to assess whether the timing of receiving explanations impacts children's generalization: is it best to give children an explanation about the mechanism before or after they get to explore the causal system? We had potential reasons for expecting either pattern of data.

Children could benefit from hearing an explanation before experiencing causal failure. Legare (2012) asked children to explain inconsistent or consistent causal events that they witnessed before exploring them. The results suggested that the explanations children generated about the inconsistent events constrained their exploration. This suggests that explanations help children build hypotheses and the following exploration serves as a space for hypothesis testing, especially in cases of causal inconsistency. Similarly, other work suggests that receiving explanations about causal structures may constrain children's exploration of causal mechanisms and promote generalization (Legare, 2014; Keil, 2006).

In Study 1, younger children spent a higher proportion of time with the toy than older children – indicating they may be taking longer to make the switch work. Hearing mechanistic explanations could constrain their hypothesis testing to a narrower mechanistic domain. This, in turn, may result in a stronger understanding of the mechanisms involved in making the causal system work – ultimately leading to children being able to apply this to novel settings.

Alternatively, it may be beneficial for children to hear a mechanistic explanation after an opportunity to explore the causal system. A wealth of research indicates that children use active exploration to learn about causal structures in the world (Bonawitz et al., 2012, Sobel and Legare, 2014, Schulz & Bonawitz, 2007). The results of Study 1 indicate that children's age does not influence their ability to interact with and reference the mechanisms in their explanations. This suggests that exploration may be sufficient for children to uncover the mechanism; the

problem may not lie in the hypothesis generation-testing relation but in children's limited ability to abstract their findings, thereby limiting their ability to generalize it to other situations.

Providing a mechanistic explanation after children's independent exploration may validate their observations and provide a clearer mechanistic framework for the causal system. This approach could motivate children to apply their findings to a novel setting, such as in the case of the similar but novel light toy.

This study served two purposes: we investigated whether mechanistic explanations improved 3- and 4-year-olds' ability to generalize learned mechanistic information to novel contexts and whether the timing of the explanations impacted a child's ability to generalize. This study was similar to Study 1 – children were given a fan toy to explore but the switch was OFF for all participants. Furthermore, we explained the mechanisms of the toy to children before or after they explored the toy. Finally, we provided them with the novel but similar light toy, where the switch was OFF. Our primary measure was whether children interacted with the switch or referenced it in response to our questions. Our primary hypothesis was that children who received an explanation would be more likely to interact with the switch on the light toy, indicating that mechanistic explanations support children's ability to generalize what they learn through their independent exploration.

Additionally, we predicted that the timing of explanations would impact children's generalization. This could be observed in two different ways: if explaining helps children's mechanistic reasoning by constraining their exploration, then children who receive an explanation before exploration may be more likely to reference the switch in the *Generalization* phase. However, if explanations help children's mechanistic reasoning by promoting their ability

to abstract acquired information on a broader level, then children who receive an explanation after exploring will be more likely to reference the switch in the *Generalization* phase

Method

Participants

We recruited 3- to 4-year-old participants from the Museum of Science, Boston. They were randomly assigned to one of three conditions: *Explain-Before*, *Explain-After*, and *Baseline* conditions. In contrast to Study 1, which included both children who did and who did not find the hidden switch during exploration, in Study 2 we were specifically interested in comparing the impact of explanation on children who found the hidden switch in the context of failure. Therefore, we recruited participants until we obtained a sample of 108 participants (36 Explain-Before, 36 Explain-After, 36 Baseline) who found the switch on the fan toy. This recruitment procedure required us to recruit a total of 206 participants; a trend we expected from Study 1 where 10 out of 17 (58.82%) children in the Failure condition found the switch on the fan toy.

Seventy-seven participants were not included in our final analyses for the following reasons: not finding the switch (12), finding the switch but not putting the toy together correctly (12), not finishing the study (20), and other reasons (parental interference, researcher error, etc; 12). An additional 21 participants were not included in our final analyses despite having a working fan toy during their exploration; these participants pushed the toy against the table so that turned the switch on surreptitiously. Since we cannot demonstrate that these participants knew about the switch on the fan toy, we did not use them in our final dataset.

Caregivers of 84 participants opted to report the child's demographic information. Out of those 84 participants, 41 were male and 37 were female. Additionally, 57 participants identified as White, 3 as Hispanic/Latine, 1 as Black, 9 as Asian, and 12 as multiracial individuals.

Additionally, 43 caregivers shared their highest level of education – 2 adults finished high school, 5 adults finished some college, 32 adults finished college, and 43 adults had an advanced degree.

Materials

We used the same toys from Study 1. However, in study 1, 29 out of 71 participants did not initially match the two parts of the toy according to our instructions. Therefore, to promote aligning the toy correctly, instead of asking children to match two thin red lines, we asked them to match the sides of the toy by color (as shown in Figure 7). Furthermore, there was no box on the fan toy further to promote younger children’s engagement with the switch. We kept the box on the light toy but modified it so that children had to match the sides instead of the lines.

Figure 7.

Stimuli used in Exploration, Explanation, and Generalization Phases

A



B



Note: A) Updated fan toy used in Exploration and Explanation phase. B) Updated light toy used in Generalization phase

Procedure

Each participant underwent five phases of the study: *Introduction* phase, *Exploration* phase, *Exploration* phase, *Distractor* phase, and *Generalization* phase.

Introduction Phase

All participants were shown a fan toy and were told about the game they could play with the toy. They were told about connecting the two parts of the toy to make it work (“*This toy works when you put these two parts together so the red parts touch*”).

Explanation Phase

Participants were told about the mechanisms of the toy (“*The toy needs power for it to work. Inside the toy, there are batteries that have power. There are also wires inside the toy, and the wires help the power go from the batteries to the fan. The toy also has a switch on it. When the switch is on, the power can move through the wires and get to the fan. When the power gets to the fan, the fan turns on.*”).

In the *Explain-Before* condition, this phase took place before the *Exploration* phase, and in the *Explain-After* condition, this phase took place after the *Exploration* phase. In the *Baseline* phase, children did not receive this explanation at all. Instead, they were asked a non-functional question about the toy (“*What colors are on this toy?*”).

Exploration Phase

All participants were given a new but identical fan toy to explore. The switch in this toy was always OFF so the fan did not turn on when the participants put the toy together.

Participants got up to 5 minutes to explore the toy and play the game.

Distractor Phase

During this phase, children were engaged in a coloring task. This ensured that children’s behavior in the *Generalization* phase was not a result of recency or primacy effects but due to the timing of the explanation.

Generalization Phase

Lastly, the experimenter handed participants the novel but similar light toy. The switch was on for this toy and it worked when children put it together correctly (according to the first causal explanation). Children did not receive instructions about the toy and were told “*I have another toy you can play with*”. After they were done exploring the toy or when a minute was over, children were asked “*What can you tell me about this toy?*” and “*How does it work?*”

Coding

Two research assistants coded each participant’s exploratory behaviors, verbal explanations, and nonverbal gestures across the *Exploration* and *Generalization* phases, following the guidelines in Study 1.

Preregistered procedure and analysis

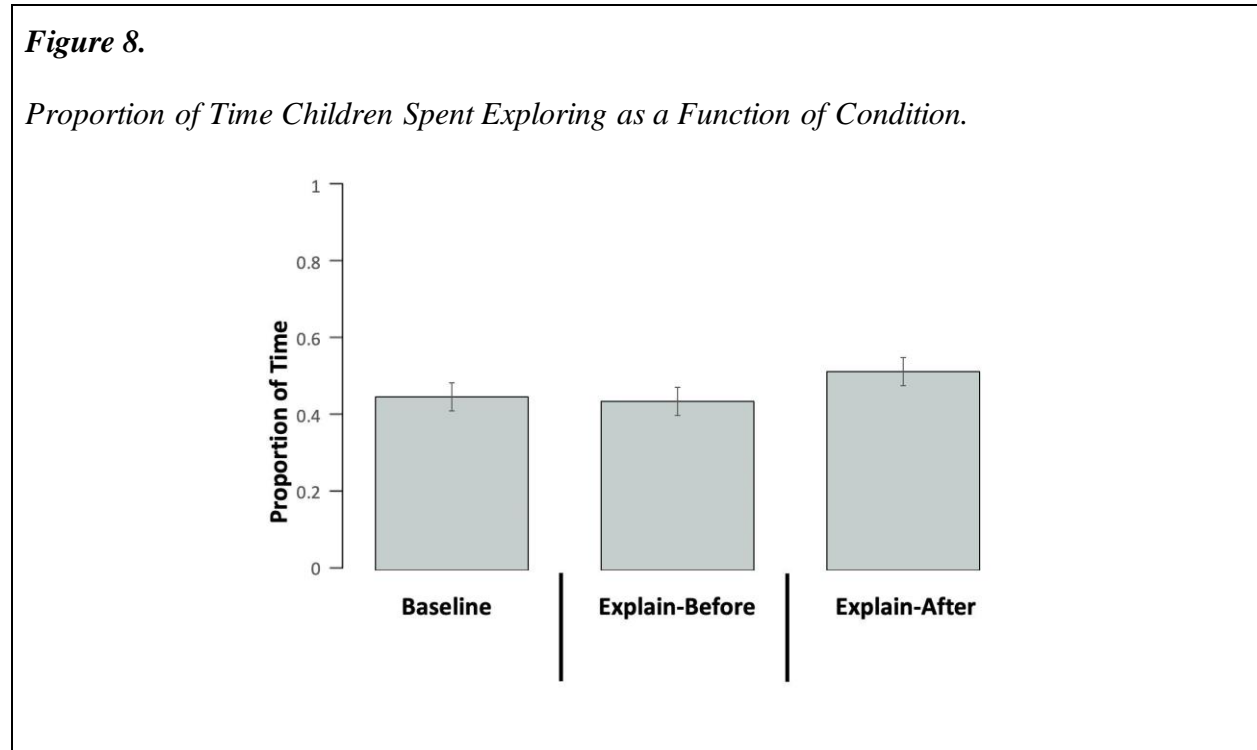
All coding measures and predictions are registered on Open Science Framework (OSF) (https://osf.io/p7sxe/?view_only=f96e3c9573b743719a382bfb2e268958).

Results

Exploration Phase

Methodologically, children in the Baseline condition and Explain-After condition had the same experience when entering the *exploration* – they were given basic instructions about the toy and told they could play with it. Children in the Explain-Before condition, however, were given a mechanistic explanation about the toy before exploration. Therefore, we chose to collapse the children in the Baseline and Explain-After condition into one group and compare them to the children in the Explain-Before condition for the *Exploration* phase analyses. The following pattern of results remains consistent when all three conditions are compared to each other. We assessed whether condition predicted the proportion of time participants spent exploring the toy by conducting a linear regression. Condition (Explore-Before vs. Explore-After / Baseline)

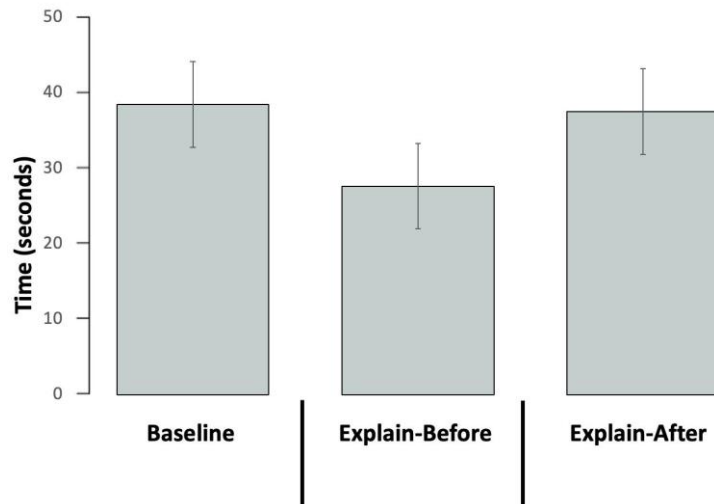
was entered as a fixed factor and it was not a significant predictor of the proportion of time children spent exploring the toy, $\chi^2(1) = 0.917, p = 0.338$ (Figure 8). Thus, there were no significant differences in the proportion of time children in the Baseline ($M = 0.45, SD = 0.20$), Explain-After ($M = 0.52, SD = 0.21$), and Explain-Before ($M = 0.44, SD = 0.25$) conditions explored the toy.



Additionally, we investigated whether condition influenced the amount of time it took children to find the switch on the fan toy. Condition (Explore-Before vs. Explore-After / Baseline) was entered as a fixed factor and it was not a significant predictor of the amount of time children took to find the switch, $\chi^2(1) = 5.06, p = 0.080$ (Figure 9). Thus, there were no significant differences in the number of seconds it took children to find the switch in the Baseline ($M = 39.00, SD = 44.70$), Explain-After ($M = 37.56, SD = 27.59$), and Explain-Before ($M = 27.64, SD = 26.60$) conditions explored the toy.

Figure 9.

Amount of time (in seconds) to find switch during Exploration as a Function of Condition



Generalization Phase

We assessed whether children generalized any mechanistic reasoning. We conducted a logistic regression on participants' likelihood to spontaneously interact with the switch (binary, 0 = No, 1 = Yes) when given the novel but similar light toy with no prompt (*"I have another toy you can play with"*); condition was entered as a dummy-coded fixed factor with Baseline condition as the reference category. Eighteen out of 36 children (50%) in the Baseline condition, 14 out of 36 children (38.89%) in the Explain-Before, and 12 out of 36 children (33.34%) in the Explain-After condition spontaneously interacted with the switch. Analyses revealed that condition was not a significant predictor of interacting with the switch, $\chi^2(2) = 2.15, p = 0.342$.


Furthermore, we investigated children's explanations when asked *"What can you tell me about the toy?"* and *"How does it work?"*. Twenty-four out of 36 children (66.67%) in the Baseline condition, 25 out of 36 children (69.44%) in the Explain-Before condition, and 18 out of 36 children (50%) in the Explain-After condition referenced the switch in their responses to

our questions during the *Generalization* phase. We conducted a logistic regression on participants' reference to the switch (binary, 0 = No, 1 = Yes) to either question with condition as a fixed factor. Contrary to our hypotheses, condition was not a significant predictor of children referencing the switch, $\chi^2(2) = 3.35, p = 0.188$.

As a final measure of generalization, we investigated whether condition predicted if children interacted with the switch when given the toy with no prompt or referenced the switch in response to our questions (“*What can you tell me about the toy?*” and “*How does it work?*”). A logistic regression was conducted on participants' reference to or interaction with the switch (binary, 0 = No, 1 = Yes); condition was entered as a fixed factor. Condition did not significantly predict children's generalization. $\chi^2(2) = 1.81, p = 0.406$ (Figure 10). Twenty-seven out of 36 children (75%) in the Baseline condition, 26 out of 36 children (72.22%) in the Explain-Before condition, and 22 out of 36 children (61.11%) in the Explain-After condition interacted with or referenced the switch in the *Generalization* phase.

Figure 10.

Proportion of Children Referencing the Switch during Exploration as a Function of Condition



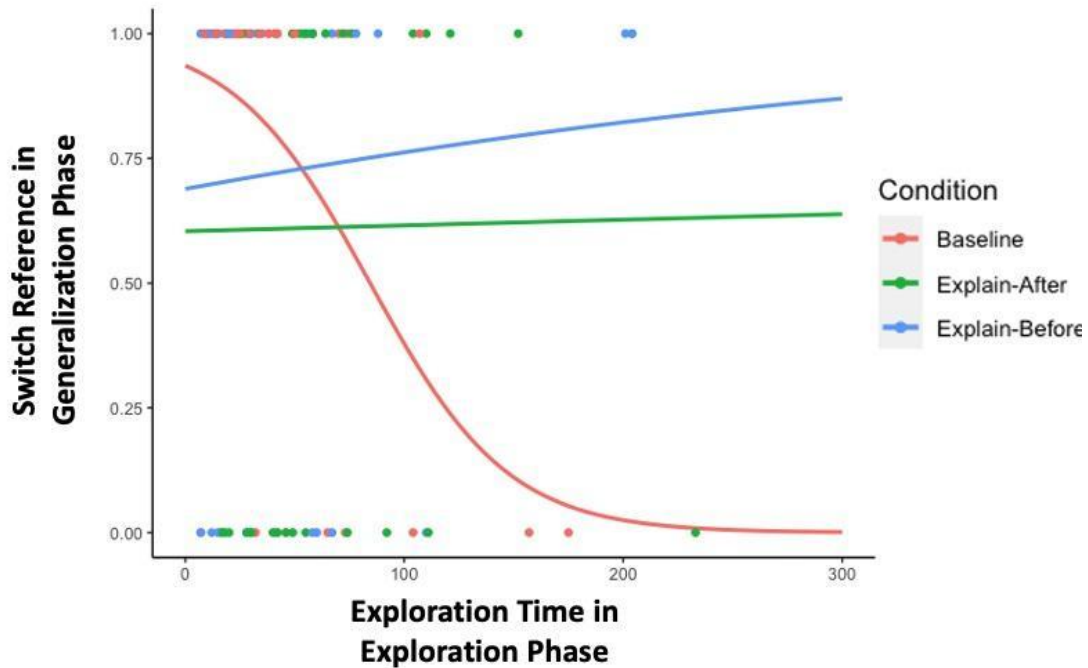
Exploratory Analyses

Similar to Study 1, we also conducted exploratory analyses to assess the relation between children's exploration behaviors and their subsequent generalization behaviors. First, we investigated whether the proportion of time children spent exploring the fan toy in the *Exploration* phase predicted whether children interacted with or referenced the light toy switch in the *Generalization* phase. We conducted a logistic regression on interaction with or reference to the switch (binary, 0 = No, 1 = Yes) with proportion of time as a continuous predictor. Proportion of exploration time did not significantly predict children's generalization, $\chi^2(1) = 1.14, p = 0.286$. Similarly, the time it took children to find the fan toy switch did not predict their reference to or interaction with the switch in the light toy, $\chi^2(1) = 1.52, p = .218$.

However, a subsequent model that included the interaction between Condition and the proportion of exploration time was significant, $\chi^2(2) = 7.46, p = .024$. Children who explored for longer in the *Exploration* phase in Explain-Before condition were more likely to interact with or reference the switch in the *Generalization* phase than children who explored for more time in the Baseline condition, $z = 2.16, p = .031$. Similarly, children who spent more time exploring the fan toy in the Explain-After condition were more likely to generalize than children who spent more time exploring in the Baseline condition, $z = 2.11, z = 0.035$ (Figure 11). However, we do not find any significant differences in children's generalization based on their exploration of the fan toy in the Explain-Before and Explain-After condition, $z = 0.28, p = 0.777$.

Figure 11.

Interaction between exploration time in Exploration phase and Condition to predict children's generalization



Finally, we compared the 3- and 4-year-olds from Study 1 to our participants in Study 2. Specifically, we compared the Failure condition of Study 1 to the Baseline condition of Study 2 since they are the most similar methodologically (both groups experience causal failure and are not provided with explanations). Due to the small sample size, we did not run any formal analyses. In Study 1, 10 3- and 4-year-olds in the Failure condition found the switch in the fan toy during the *Exploration* phase. In Study 2, 36 3- and 4-year-old children in the Baseline condition found the switch in the *Exploration* phase. However, only 4 children interacted with or referenced the switch in the *Generalization* phase of Study 1 (40%) whereas 27 out of 36 children interacted with or referenced the switch in Study 2 (75%).

Discussion

The current study investigated the role of explanations in children's generalization of acquired mechanistic knowledge. We gave 3- to 4-year-old children a toy (a fan) and instructions on how to use it to play a goal-directed game. Children were given basic instructions about putting the toy together but were not told about a partially hidden mechanism (i.e., a switch) set to the OFF position, leading to causal failure. The study had three conditions: no explanation about the fan toy was provided (Baseline condition), an explanation was provided before encountering causal failure (Explain-Before condition), and an explanation was provided after experiencing causal failure (Explain-After condition). Subsequently, children were presented with a novel but similar light toy where the switch was ON to assess if children would spontaneously apply their learned knowledge to this new situation.

Our primary hypothesis was that providing explanations would enhance children's ability to generalize learned mechanistic information, with children in the Explain-Before and Explain-After conditions generalizing at higher rates than children in the Baseline condition. Additionally, we aimed to discern the role of explanations by manipulating the time at which they were provided. We posited that explanations would either help children frame their hypotheses to test in their subsequent exploration (Explain-Before condition) or explanations would aid in abstracting information acquired from their previous exploration (Explain-After condition).

Study 1 of this project identified an effect of age on children's ability to generalize learned mechanistic information – specifically, only 36.37% of the younger children mentioned the switch in the light toy, compared to 75% of the older children. Based on these results, Study 2 aimed to bridge the gap between older and younger children by building on research indicating that explanations enhance children's learning and ability to generalize mechanistic knowledge.

Contrary to our predictions, providing explanations did not significantly impact children's ability to generalize mechanistic knowledge – 75% of children in the Baseline condition, 72.22% in the Explain-Before condition, and 59.46% of children in the Explain-After condition referred to or interacted with the switch in the light toy. However, children in these conditions outperformed 3- to 4-year-olds in Study 1 where only 40% of children who experienced causal failure referred to or interacted with the switch in the light toy.

Furthermore, in our exploratory analyses for Study 2, we identified an interaction between condition and the duration of time children spend exploring the fan toy. Specifically, children who spent more time exploring the fan toy in the Baseline condition during the *Exploration* phase were less likely to find the light switch in the *Generalization* phase compared to children who spent a longer time exploring the fan toy in the Explain-Before or Explain-After conditions. This suggests that extended exploration without receiving an explanation limits children's ability to generalize learned mechanistic information, providing us with some evidence of when explanations may support children's generalization.

The discrepancy between children in the Baseline condition of Study 2 and 3- to 4-year-olds in the Failure condition of Study 1 prompts us to consider the methodological difference between studies. In Study 1, children were asked questions about the functioning of the fan toy (“*What can you tell me about the toy?*” and “*How does it work?*”), whereas in the Baseline condition of Study 2, children were asked a perceptual question about the fan toy (“*What colors are on this toy?*”). Additionally, in Study 1, children immediately transitioned from exploring and explaining the fan toy to receiving the novel but similar light toy. In contrast, in Study 2, after children's exploration and explanation, there was a 5-minute distractor task prior to receiving the light toy. Lastly, in Study 1, the experimenter handed the light toy to children saying they did not know how the toy worked and needed help figuring it out before proceeding

to ask them questions about the toy. In Study 2, however, the experimenter simply stated that they had another toy for children to play with, after which they asked them questions about the toy. These differences in prompting children to explain the fan toy and asking them for help in figuring out how the light toy works might suggest that children in the Failure condition of Study 1 would be more likely to generalize information than those in the Baseline condition of Study 2. However, our results show the opposite trend.

While these methodological differences in studies do not elucidate why there are differences in the generalization rates of 3- and 4-year-olds in the Failure condition of Study 1 and children in the Baseline condition of Study 2, one plausible explanation is the inclusion of the *Distraction* phase. The inclusion of this phase could allow children the opportunity to consolidate information gathered in the *Exploration* phase, allowing them to develop a clearer framework for how these toys worked. Additionally, in Study 2, both toys were changed so that instead of asking children to match thin red lines, we asked them to match the sides of the toy by color. The fan toy in Study 2 was also modified so that the switch was more accessible to children than it was in Study 1. While no changes were made to the accessibility of the switch in the light toy, these structural changes may have prompted children's ability to generalize the knowledge of the switch from the fan toy to the light toy. However, further research, including direct comparisons of children's generalization abilities with and without a *Distraction* phase, is necessary to develop a fuller picture of how children apply acquired mechanistic knowledge to novel situations.

General Discussion

This project comprised two studies aimed at understanding the contexts that promote children's mechanistic reasoning. Study 1 investigated the impact of causal failure (as opposed to causal success) in 3- to 6-year-old children's exploration, explanation, and generalization of

mechanisms. As predicted, the results of Study 1 indicated that causal failure is a context that promotes children's exploration and explanation of the relevant mechanisms. We also found that although older children spent proportionally less time overall playing with the toy, they were still more likely to interact with the hidden switch than younger children. Finally, we found that 5- and 6-year-old children were more likely than 3- and 4-year-old children to generalize their mechanistic reasoning.

Study 2 was built on the basis of the results from Study 1 and investigated the role of explanations in 3- to 4-year-old children's generalization from causal failure. This study aimed to promote younger children's generalization by providing them with mechanistic explanations. Furthermore, the study investigated the role of explanations by manipulating the time at which they were provided (before or after children's exploration of causal failure). Contrary to our predictions, we did not find an impact of explanations on children's generalization or any differences based on the timing of explanations. Intriguingly, we found that children who did not receive explanations in Study 2 generalized at higher rates than those in Study 1 (who also did not receive an explanation). This unexpected discrepancy may be attributed to the methodological differences in the two studies. Further research is needed to develop a comprehensive understanding of children's generalization of mechanistic information and the role explanations may play in that. Finally, we found that children who explore causal failure for a longer period generalize at higher rates when they receive an explanation (before or after the exploration) than when they do not receive an explanation.

Role of failure in mechanistic reasoning

Prior research has shown that children's causal exploration is enhanced in cases of ambiguity. For example, when potential causes present confounded evidence, children overcome

novelty preferences and explore this confounded evidence more, and in some cases engage in behaviors that could disambiguate that evidence (Cook et al., 2011; Schulz & Bonawitz, 2007). Children also imitate causal actions more reliably when they deterministically produce outcomes compared to when they probabilistically produce outcomes (Schulz, Hoopell, et al., 2008). Finally, 3- to 7-year-old children explore more – and 6- to 7-year-old children are more likely to appeal to potential mechanisms – when they are faced with evidence inconsistent with their prior beliefs (Bonawitz et al., 2012). Given failure is implicit in many of these paradigms, the current project provides a clearer view of the role of failure in children’s mechanistic reasoning across two studies.

Two age-related differences emerged from Study 1. First, while younger children engaged in more exploration than older children, older children were more likely to discover the hidden switch relative to younger children. This finding speaks partly to the means by which contexts of failure support (or perhaps do not support) mechanistic reasoning.

Second, we found an age-related change in children’s generalization of mechanistic reasoning. Older children were more likely to discover a hidden switch in a new, but perceptually similar, functioning toy and more likely to explain how this new toy worked by referencing the hidden switch. Critically, given children in all conditions succeeded during the *Generalization* phase, the effect of failure contexts on their mechanistic reasoning in older children suggests that failure impacted children’s generalization of their prior learning, rather than new learning of another toy that also failed to work.

Relation between exploration, explanation, and generalization

The current project included several measures of children’s mechanistic reasoning – their exploratory behaviors, their explanations, as well as how they generalized these behaviors to new contexts. How are these different aspects of causal exploratory play related? From Study 1, we

know that one possibility is that these abilities are linearly related. That is, exploring more may allow children to explain causal relations better, and enhanced exploration and explanation may support generalizing this knowledge to new contexts. If so, then the current findings may suggest that failure promotes both children's exploration and explanation, or instead promotes only their exploration. The current study cannot disentangle these alternatives. Given contexts of failure had such strong effects on children's exploration, there were very few children in the Failure condition who did not find the hidden switch or children in the Success condition who did find it. Thus, we could not assess whether children in the Failure condition who found the switch would be more likely to explain mechanisms compared to children in the Success condition who also found the switch. Further investigations using a paradigm in which children in Success conditions spontaneously discover mechanisms at higher rates than the current study would shed additional light on this question.

Study 2 delved deeper into understanding the relation between exploration and generalization in younger ages by providing children with explanations. All children experienced causal failure; however, they were either given no explanations about the causal system or given mechanistic explanations about the system before or after they explored the causal failure. We aimed to determine whether providing mechanistic explanations would enhance children's generalization abilities, and if the timing of explanations would play a role.

While our hypotheses were contradicted and we did not find an impact of explanations or the timing of explanations, our exploratory analyses revealed an interaction between children's exploration time and whether or not they were given an explanation. In Study 1, we observed that simply exploring the toy longer did not enhance children's mechanistic reasoning, suggesting that exploration time alone cannot predict children's generalization abilities. In Study 2, we added nuance to this finding: duration of exploration did not alone predict children's

ability to generalize mechanistic information to novel contexts. However, longer exploration without explanations led to lower rates of generalization compared to when explanations were provided (either before or after exploration). These results could potentially contribute to the research about children's persistence abilities. Previous research suggests that children tend to persist more when observing adults persist, when they receive praise, and when adults intervene less (Leonard, 2021). However, persistence without explanation, as occurred in the Baseline Condition in Study 2, might not improve children's ability to apply what they learned to novel situations. Future studies should further investigate the relation between exploration time, providing explanations, and generalizing information.

Additionally, we could investigate the types of explanations that facilitate generalization. We know that children's ability to generalize mechanistic information increases when they receive mechanistic (compared to non-mechanistic) explanations (Kurkul et al, 2021). However, the explanations provided in this work comprised both failure-oriented and success-oriented language. Future work could investigate whether failure-oriented mechanistic explanations affect children's exploration, explanation, and generalization differently than success-oriented mechanistic explanations.

Study 2 aimed to better understand how receiving mechanistic explanations supports children's mechanistic reasoning in the context of causal failure. We proposed two possibilities: explanations could either constrain the range of children's subsequent explanations to a relevant domain, or explanations might be most effective after children's exploration, allowing them to abstract the knowledge they obtained from exploration. Our stimuli lacked a variety of features for children to explore – apart from assembling the toy and turning the switch ON or OFF, children could not explore any mechanistic features of the toy. Future studies could better

understand the role of explanations in exploration and generalization by providing children with a causal system that has more facets to be explored.

Finally, the current project leaves open questions concerning how we should conceptualize the role of failure in supporting children's mechanistic reasoning. Is failure another instance of ambiguity or surprise, both contexts which may encompass contexts of failure? Or does failure play a specific role in prompting children's mechanistic reasoning? Several distinct literatures have also proposed that surprise (Stahl & Feigenson, 2015) and uncertainty (Gopnik et al., 2001) both play a critical role in children's causal learning. Although all cases of failure may involve uncertainty and surprise, not all cases of uncertainty and surprise involve the failure: children may be uncertain about which of two potential actions caused a particular outcome or be surprised by the function of a new toy. Future studies should aim to characterize more precisely the precise role of failure in children's mechanistic reasoning and how it relates conceptually to other contexts shown to facilitate children's causal learning.

Conclusion

In conclusion, this project provides evidence that causal failure, but not success, prompts children's exploration and explanation of the relevant mechanisms. Additionally, we know that older preschool-aged children are more likely to generalize mechanistic learning to devices with similar causal structures. For younger preschoolers, we do not find an impact of mechanistic explanations (or their timing) on their ability to generalize mechanistic learning. Nonetheless, we observe that children who spend more time exploring are more likely to generalize mechanistic learning when given an explanation. Overall, our findings suggest that causal failure may enhance mechanistic reasoning, and future studies should aim to clarify the role of explanations in children's generalization.

References

- Gweon, H., & Schulz, L. (2011). 16-month-olds rationally infer causes of failed actions. *Science*, *332*(6037), 1524-1524.
- 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. <https://doi.org/10.1111/cdev.12613>
- Bonawitz, E. B., Ferranti, D., Saxe, R., Gopnik, A., Meltzoff, A. N., Woodward, J., & Schulz, L. E. (2010). Just do it? Investigating the gap between prediction and action in toddlers' causal inferences. *Cognition*, *115*(1), 104–117. <https://doi.org/10.1016/j.cognition.2009.12.001>
- Bonawitz, E. B., van Schijndel, T. J. P., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, *64*(4), 215–234. <https://doi.org/10.1016/j.cogpsych.2011.12.002>
- Buchanan, D., & Sobel, D. (2011). Children posit hidden causes to explain causal variability. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 33, No. 33).
- Buchanan, D. W., & Sobel, D. M. (2011). Mechanism-Based Causal Reasoning in Young Children. *Child Development*, *82*(6), 2053–2066. <https://doi.org/10.1111/j.1467-8624.2011.01646.x>
- Buchsbaum, D., Bridgers, S., Skolnick Weisberg, D., & Gopnik, A. (2012). The power of possibility: Causal learning, counterfactual reasoning, and pretend play. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *367*(1599), 2202–2212. <https://doi.org/10.1098/rstb.2012.0122>

Callanan, M. A., & Oakes, L. M. (1992). Preschoolers' questions and parents' explanations: Causal thinking in everyday activity. *Cognitive Development*, 7(2), 213–233.

[https://doi.org/10.1016/0885-2014\(92\)90012-G](https://doi.org/10.1016/0885-2014(92)90012-G)

Erb, C. D., Buchanan, D. W., & Sobel, D. M. (2013). Children's developing understanding of the relation between variable causal efficacy and mechanistic complexity. *Cognition*, 129(3), 494-500.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A Theory of Causal Learning in Children: Causal Maps and Bayes Nets. *Psychological Review*, 111(1), 3–32. <https://doi.org/10.1037/0033-295X.111.1.3>

Gopnik, A., & Schulz, L. (2004). Mechanisms of theory formation in young children. *Trends in Cognitive Sciences*, 8(8), 371–377. <https://doi.org/10.1016/j.tics.2004.06.005>

Gopnik, A., & Sobel, D. M. (2000). Detecting Blickets: How Young Children Use Information about Novel Causal Powers in Categorization and Induction. *Child Development*, 71(5), 1205–1222. <https://doi.org/10.1111/1467-8624.00224>

Gopnik, A., Sobel, D. M., Schulz, L. E., & Glymour, C. (2001). Causal learning mechanisms in very young children: Two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Developmental Psychology*, 37(5), 620–629.

<https://doi.org/10.1037/0012-1649.37.5.620>

Gweon, H., & Schulz, L. E. (n.d.). *Stretching to learn: Ambiguous evidence and variability in preschoolers' exploratory play*.

Keil, F. C. (2006). Explanation and Understanding. *Annual Review of Psychology*, 57, 227–254.

<https://doi.org/10.1146/annurev.psych.57.102904.190100>

- Kelemen, D. (2019). The Magic of Mechanism: Explanation-Based Instruction on Counterintuitive Concepts in Early Childhood. *Perspectives on Psychological Science*, *14*(4), 510–522. <https://doi.org/10.1177/1745691619827011>
- Kurkul, K. E., Castine, E., Leech, K., & Corriveau, K. H. (2021). How does a switch work? The relation between adult mechanistic language and children’s learning. *Journal of Applied Developmental Psychology*, *72*, 101221. <https://doi.org/10.1016/j.appdev.2020.101221>
- Leech, K. A., Haber, A. S., Jalkh, Y., & Corriveau, K. H. (2020). Embedding Scientific Explanations Into Storybooks Impacts Children’s Scientific Discourse and Learning. *Frontiers in Psychology*, *11*. <https://doi.org/10.3389/fpsyg.2020.01016>
- Legare, C. H. (2012). Exploring Explanation: Explaining Inconsistent Evidence Informs Exploratory, Hypothesis-Testing Behavior in Young Children. *Child Development*, *83*(1), 173–185. <https://doi.org/10.1111/j.1467-8624.2011.01691.x>
- Legare, C. H. (2014). The Contributions of Explanation and Exploration to Children’s Scientific Reasoning. *Child Development Perspectives*, *8*(2), 101–106. <https://doi.org/10.1111/cdep.12070>
- 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. <https://doi.org/10.1111/cdev.13246>
- Lombrozo, T., Bonawitz, E. B., & Scalise, N. R. (2018). Young children’s learning and generalization of teleological and mechanistic explanations. *Journal of Cognition and Development*, *19*(2), 220–232. <https://doi.org/10.1080/15248372.2018.1427099>
- Mills, C. M., Danovitch, Judith. H., Mugambi, V. N., Sands, K. R., & Pattisapu Fox, C. (2022). “Why do dogs pant?”: Characteristics of parental explanations about science predict

children's knowledge. *Child Development*, 93(2), 326–340.

<https://doi.org/10.1111/cdev.13681>

Muentener, P., & Bonawitz, E. (2018). *The development of causal reasoning*. OSF Preprints.

<https://doi.org/10.31219/osf.io/r8h9p>

Muentener, P., & Schulz, L. (2014). Toddlers infer unobserved causes for spontaneous events.

Frontiers in Psychology, 5. <https://doi.org/10.3389/fpsyg.2014.01496>

Nazzi, T., & Gopnik, A. (2000). A shift in children's use of perceptual and causal cues to categorization. *Developmental Science*, 3(4), 389–396. <https://doi.org/10.1111/1467-7687.00133>

Russ, R. S., Scherr, R. E., Hammer, D., & Mikeska, J. (2008). Recognizing mechanistic reasoning in student scientific inquiry: A framework for discourse analysis developed from philosophy of science. *Science Education*, 92(3), 499–525.

<https://doi.org/10.1002/sce.20264>

Schulz, L. (2012). The origins of inquiry: Inductive inference and exploration in early childhood.

Trends in Cognitive Sciences, 16(7), 382–389. <https://doi.org/10.1016/j.tics.2012.06.004>

Schulz, L., & Bonawitz, E. (2007). Serious Fun: Preschoolers Engage in More Exploratory Play When Evidence Is Confounded. *Developmental Psychology*, 43, 1045–1050.

<https://doi.org/10.1037/0012-1649.43.4.1045>

Schulz, L. E., Goodman, N. D., Tenenbaum, J. B., & Jenkins, A. C. (2008). Going beyond the evidence: Abstract laws and preschoolers' responses to anomalous data. *Cognition*, 109(2), 211–223. <https://doi.org/10.1016/j.cognition.2008.07.017>

Schulz, L. E., & Gopnik, A. (2004). Causal learning across domains. *Developmental Psychology*, 40(2), 162–176. <https://doi.org/10.1037/0012-1649.40.2.162>

Schulz, L. E., & Sommerville, J. (2006). God Does Not Play Dice: Causal Determinism and Preschoolers' Causal Inferences. *Child Development*, 77(2), 427–442.

<https://doi.org/10.1111/j.1467-8624.2006.00880.x>

Sobel, D. M., & Legare, C. H. (2014). Causal learning in children. *WIREs Cognitive Science*, 5(4), 413–427. <https://doi.org/10.1002/wcs.1291>

Sobel, D. M., & Stricker, L. W. (2022). Messaging Matters: Order of Experience with Messaging at a STEM-Based Museum Exhibit Influences Children's Engagement with Challenging Tasks. *Visitor Studies*, 25(1), 104–125. <https://doi.org/10.1080/10645578.2021.2015948>

Stephan, S., & Waldmann, M. R. (2022). The role of mechanism knowledge in singular causation judgments. *Cognition*, 218, 104924.

<https://doi.org/10.1016/j.cognition.2021.104924>

Walker, C. M., Lombrozo, T., Legare, C. H., & Gopnik, A. (2014). Explaining prompts children to privilege inductively rich properties. *Cognition*, 133(2), 343–357.

<https://doi.org/10.1016/j.cognition.2014.07.008>

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.

<https://doi.org/10.1111/cdev.12590>

Willard, A. K., Busch, J. T., Cullum, K. A., Letourneau, S. M., Sobel, D. M., Callanan, M., & Legare, C. H. (2019). Explain this, explore that: A study of parent–child interaction in a children's museum. *Child Development*, 90(5), e598-e617.