Technische Universität München
TUM School of Social Sciences and Technology

# Investigating preschoolers’ active and ecological learning with novel non-verbal paradigms

Nora Swaboda

Vollständiger Abdruck der von TUM School of Social Sciences and Technology der Technischen Universität München zur Erlangung einer

Doktorin der Philosophie (Dr. phil)

genehmigten Dissertation.

Vorsitz: Prof. Dr. Anna Keune

Prüfer*innen der Dissertation:

1. Prof. Dr. Azzurra Ruggeri
2. Prof. Dr. Björn Meder

Die Dissertation wurde am 30.09.2022 bei der Technischen Universität München eingereicht und durch die TUM School of Social Sciences and Technology am 06.04.2023 angenommen.

# Acknowledgements

When I set out on this doctoral journey, I did not even dare to imagine the moment when I would finally say: I did it! Admittedly, juggling parenthood and being a PhD student in the midst of a previously unseen pandemic has not always been an easy feat. But the day has finally come, and I would like to express my deepest gratitude to all the people who have encouraged and accompanied me along the way.

First and foremost, I would like to thank my supervisor, Prof Dr Azzurra Ruggeri. From the first day until the last, she has provided me with invaluable support and guidance, as well as great amounts of understanding and patience. Her enthusiasm and knowledge in our field have been a great source of inspiration and motivation. I would also like to thank her for her high expectations and trust in my abilities, as, by always pushing a tiny bit further, she has allowed me to grow both personally and professionally.

My gratitude also goes to Prof Dr Björn Meder for his additional guidance, his thoughtful and detailed feedback and his help with more complex analysis and programming issues. I would also like to thank him for joining the examination committee. In this context, I would also like to express my gratefulness to Prof Dr Anna Keune for agreeing to chairing the committee.

Furthermore, I would like to thank all my fellow PhD students and colleagues of the iSearch lab for always having an open ear, for providing useful insights and advice on my projects, and for making working in the lab such a warm and welcoming experience. My special thanks also go to my many research assistants who, over the years, have greatly helped me with preparing study materials, data collection and coding. Without their assistance, I would not be where I am today.

Finally, I am deeply grateful for my partner Damien for believing in and encouraging me, but for also reminding me to take a step back and put things into perspective when necessary. Last but not least, I would like to thank my daughter Colette who, regardless of the sleepless nights she has given me, reminds me every day of the incredible wonder children truly are.

# Abstract

Play has been considered essential for learning by developmental scientists and educators alike for more than a hundred years, and a significant body of work demonstrates that children purposefully engage with their environment from a very young age. However, evidence about the efficiency of children's self-guided exploration is still mixed, with some studies showing that even preschoolers are effective active learners, highly attuned to learning opportunities and capable of designing informative interventions, and others suggesting that efficient exploration only emerges during the primary school years. This dissertation investigates the efficiency and adaptiveness of the active learning strategies of preschool-aged children using a range of novel non-verbal exploration tasks. Its results show that, under the right conditions, children as young as 3 search for information efficiently and adaptively, identifying and implementing those exploratory strategies best suited to provide them with the sought information given the current learning context. At the same time, this dissertation also outlines the limits of children's search efficiency, demonstrating that general task-demands, such as task complexity or reliance on verbal skills, directly affect the efficiency of children's exploratory actions, and that children's self-guided exploration is most efficient when the search space is constrained. Together, these results add to the growing body of literature demonstrating that young children are skilled active learners and highlight the need for age-appropriate paradigms to gain a better picture of their emerging information-search abilities. By providing insights into the cognitive mechanisms underlying children's learning and the conditions under which they explore most efficiently, the results of this dissertation are also relevant for the educational debate of whether play constitutes a suitable means for educating young children, and point towards free exploration within structured environments as a potential educational approach that combines benefits of play with a focus on enhancing learning outcomes.

# Zusammenfassung

Seit mehr als einhundert Jahren betrachten Entwicklungspsychologen und Pädagogen Spiel als essenziell für kindliche Lernprozesse, und zahlreiche Studien belegen, dass sich Kinder bereits in sehr jungem Alter gezielt mit ihrer Umwelt auseinandersetzen. Dennoch ist die Datenlage bezüglich der Wirksamkeit kindlicher Lernstrategien weiterhin uneindeutig. Einerseits zeigen Studien, dass Kinder bereits im Vorschulalter effektive aktive Lernstrategien besitzen, Lerngelegenheiten zuverlässig erkennen und in der Lage sind aufschlussreiche und informative Experimente zu entwerfen. Andere Studien jedoch deuten darauf hin, dass kindliche Lernstrategien erst im Grundschulalter wirklich effizient werden. Dieses Promotionsprojekt untersucht anhand einer Reihe neu entwickelter, nonverbaler Suchaufgaben, wie effizient und anpassungsfähig die aktiven Lernstrategien von Vorschulkindern sind. Die Ergebnisse der vorgestellten Studien zeigen, dass Kinder im Alter von 3 Jahren unter den richtigen Bedingungen effizient nach Informationen suchen und flexibel diejenigen Lernstrategien anwenden, welche in der gegebenen Lernsituation am besten geeignet sind, die gewünschten Informationen zu erschließen. Gleichzeitig zeigt dieses Promotionsprojekt jedoch auch die Grenzen kindlicher Sucheffizienz auf. Insbesondere weisen die Ergebnisse darauf hin, dass sich allgemeine Aspekte einer Aufgabe, wie z.B. deren Abhängigkeit von verbalen Fähigkeiten oder deren Komplexität, direkt auf die Effizienz kindlicher Lernstrategien auswirken und dass Kinder am besten nach Informationen suchen, wenn der durch die Aufgabe gegebene Suchraum begrenz ist. Zusammengenommen ergänzen die Ergebnisse dieses Promotionsprojekts frühere Studien, die zeigen, dass Kinder bereits in jungem Alter wirksame aktive Lernstrategien besitzen und unterstreichen, dass ein genaues Bild sich entwickelnder kindlicher Kompetenzen nur dann gezeichnet werden kann, wenn diese mit altersgerechten Paradigmen untersucht werden. Außerdem sind die Ergebnisse dieser Promotion relevant für die bildungspolitische Debatte, ob Spiel ein geeignetes Mittel für die Bildung kleiner Kinder darstellt, da sie Aufschluss geben über die kognitiven Mechanismen, die kindlichem Lernen zugrunde liegen sowie über die Bedingungen, unter denen Kinder am wirksamsten erkunden. Insbesondere sprechen die Ergebnisse für freies Spiel in strukturierten Lernumgebungen als einen möglichen pädagogischen Ansatz, der Vorteile des freien Spiels mit einem verstärkten Fokus auf das Erreichen von objektiven Lernzielen vereinbart.

# Contents

# 1. Introduction

Early childhood is increasingly being recognised as a crucial period for learning and development. Attending early childhood education is associated with better outcomes for children's academic and socio-emotional development, with especially large benefits for children from economically disadvantaged backgrounds (e.g., Bai et al., 2020; Magnuson et al., 2007; McCoy et al., 2017; Yoshikawa et al., 2013). Accordingly, governments and policy makers are increasingly trying to improve early childhood education. However, there are substantial differences in the ways early childhood education is conceptualised and implemented across countries. In Germany for example, play- and situation-oriented approaches, which focus on providing children with opportunities to explore their surroundings and interests on their own, are the norm (e.g., Faas et al., 2017). In contrast, the prevailing view in the US is that children learn best through formal teaching and direct instruction, and periods of free play are progressively being replaced by more structured teaching sessions and standardised testing documenting children's academic achievements (e.g., Brown, 2021; Ilgaz et al., 2018). At the same time, critical voices are being raised in both countries: in Germany, where more structured curricula emphasising the development of key competencies in language, maths and early science skills have only been implemented during the early 2000s, there is concern that purely play-based approaches do not prepare children well enough for the demands of formal education (e.g., Faas et al., 2017; Roßbach, 2008). In the US on the other hand, critics point out that purely achievement-oriented early childhood practices put too much pressure on young children and do not allow them enough time for free play and self-guided exploration (e.g., Hirsh-Pasek et al., 2020; Nicolopoulou, 2010; Sahlberg & Doyle, 2019; Yogman et al., 2018). Despite apparently opposing positions, these two discourses share a common core, namely the question of whether children learn better (or at all) during play or through formal instruction. The idea that play is essential for learning has been present in education and developmental psychology for more than a century, and is supported by an impressive body of empirical work demonstrating that children engage with their environment in meaningful ways from very early on. However, evidence about the efficiency of children's self-guided exploration is still mixed, with some studies demonstrating that even toddlers and preschoolers possess effective inquiry skills, and others suggesting that efficient exploration only emerges during the primary school years.

This dissertation explores children's active learning from a cognitive developmental perspective. Introducing a range of novel age-appropriate non-verbal behavioural paradigms, it investigates whether preschoolers' active learning strategies are *efficient* (i.e., optimally provide children with the sought information), and whether they are *adaptive*, that is, whether children flexibly tailor their exploratory strategies to the given learning context. In answering these questions, this projects aims to provide a better understanding of the cognitive processes underlying children's active learning, as well as to inform the debate about the educational suitability of free play in early childhood.

## 1.1. Learning through play: Theoretical approaches

The idea that play supports learning has been endorsed by developmental scientists and educators alike for more than a hundred years (e.g., Berlyne, 1969; Bruner et al., 1976; Fröbel, 1839; Montessori, 1912; Piaget, 1954; Vygotsky & Wollock, 1997). Maria Montessori for instance considered play "the child's work" (1967, p.180), during which children explore their surroundings following their own interests and thereby foster their understanding of the world around them. Similarly, Friedrich Fröbel, founder of the first pedagogical childcare centres in $19^{th}$ century Germany, saw play both as a means of self-expression and of understanding the world, and his child-centred pedagogy, encouraging children to actively engage in purposeful play (Fröbel, 1885), still influences German early childhood education today.

In psychology, active involvement with their physical and social environment has also been considered crucial for children's learning since the earliest developmental theorists (e.g., Dewey, 1986; Piaget, 1954; Vygotsky & Wollock, 1997). Jean Piaget for instance proposed that children actively *construct* knowledge through a *"constant process of going back and forth between the person and the environment"* (Piaget, 1929), where learning occurs when the environment does not behave in ways consistent with children's prior knowledge and expectations of how the environment should behave. This mismatch between what is expected and what is observed generates a state of cognitive disequilibrium, which children are motivated to resolve either by re-interpreting new information to fit into their existing conceptual structures (assimilation), or by adapting existing conceptual structures to more accurately reflect new observations (accommodation).

Piaget's theory remains the starting point for most of today's work in cognitive developmental psychology. Over the past decades, his constructivist account of learning has been reformulated and refined, for instance by conceptualising learning as a process of *theory revision*, during which children, not unlike scientists, generate intuitive theories about the world and revise those theories when encountering new or conflicting evidence (e.g., Carey, 1985; Gopnik & Meltzoff, 1997; Gopnik & Wellman, 1992; L. Schulz, 2012; Wellman & Gelman, 1992; Xu, 2019). A major drawback of these purely cognitive accounts however is that they do not provide details of how the processes of knowledge construction and theory revision might work, which makes it difficult to derive empirically testable predictions from them. More recently, these models have therefore been grounded in computational models providing tentative accounts of how theory revision, and cognitive development more generally, might take place. Most prominently, Bayesian probabilistic models assume that children derive intuitive theories about the world from statistical patterns and regularities they observe in the environment, and adapt existing theories when encountering new conflicting evidence (e.g., Gopnik & Wellman, 2012; Xu & Griffiths, 2011; Xu & Kushnir, 2013). Importantly, theory revision and adaptation are assumed to proceed according to the computational principles of Bayesian inference, which describes and quantifies how much new evidence is needed to override or change an existing theory or belief, given the strength of the initially held belief. By providing clearly quantifiable mechanisms of how and when theory revision should occur, these computational models make clear empirical predictions that can be tested in behavioural studies. Accordingly, the past decades have seen a great surge in empirical investigations on children's learning, and there is now robust evidence supporting the idea that children learn through meaningful interaction with their environment from a very young age, and possibly from the moment they are born.

## 1.2. Constructing intuitive theories: Learning from observable statistical patterns

During the first months of life, infants primarily explore their surroundings visually or orally (e.g., Butterworth & Hopkins, 1988; Rochat, 1983, 1987, 1989). Studies investigating active learning in such young children typically present them with different stimuli, scenes, or events and record the direction and duration of their eye gaze in order to infer their degree of interest in the presented materials. Collectively, these looking-time studies have yielded rich insights into infants' conceptual representations, as well as into the cognitive processes that mediate learning (for reviews, see Aslin, 2007; Oakes, 2010, 2017). In line with the idea that children generate intuitive theories about the world from the limited statistical evidence they observe, these studies have demonstrated that infants are highly sensitive towards environmental patterns and probabilities. For example, 8-

months old are able to detect and later recognise frequently co-occurring shape- or sound-pairs in streams of visual or auditory input - an ability that is thought to be the basis for word- and object-learning (e.g., Fiser & Aslin, 2002; Kirkham et al., 2002; Saffran et al., 1996).

What's more, looking time-studies have also shown that infants are capable of drawing powerful inductive inferences based on small amounts of data ( e.g., Denison, Reed et al., 2013; Denison & Xu, 2014, 2019; Xu & Garcia, 2008). In one study, 8-months olds first watched how five red and white ping-pong balls were drawn randomly from a container, and were subsequently shown the contents of that container (i.e., the population of red and white balls). When the composition of the population was revealed, infants looked significantly longer at the display when its composition (e.g., mainly white balls) did not match the composition of the previously drawn sample (e.g., mainly red balls) than when their compositions were consistent, indicating that they had inferred the composition of the population from the limited evidence they had observed (Xu & Garcia, 2008). Other studies have also found that infants look longer when an agent who initially selected one object over another multiple times subsequently selects the previously non-chosen object (Duh & Wang, 2019; Kushnir et al., 2010; Wellman et al., 2016), and that both infants and preschoolers infer cause-and-effect relationships from how frequently different objects activate a machine (e.g., Denison, Bonawitz et al., 2013; Gopnik et al., 2001; Waismeyer et al., 2015). Together, these results demonstrate that children rely on environmental patterns of evidence to draw powerful inferences about the functioning of the physical and social worlds from a very early age.

Over the preschool years, these probabilistic reasoning skills become even more sophisticated as children begin to integrate prior probabilities with feedback and subsequent evidence (Girotto & Gonzalez, 2008; Gonzalez & Girotto, 2011), or use evidence from interventions to disambiguate causal relationships within a system (L. Schulz, Gopnik et al., 2007). There is also evidence that preschoolers revise their theories in accordance with the principles of Bayesian inference (e.g., Eaves Jr & Shafto, 2012; L. Schulz, Bonawitz et al., 2007): Schulz and colleagues (2007) showed 3- to 5-year-olds story books where, over the course of a week, a bunny sometimes developed a tummy ache after certain events and children were asked to infer what was causing the bunny's tummy ache. Building on findings that children of that age tend to have a strong prior belief that physiological states (e.g., a tummy ache) should have physiological, rather than psychological, causes (Notaro et al., 2002), there were two types of possible causes for the bunny's tummy ache: a physiological cause (e.g.., ingesting a certain type of food), or a psychosomatic cause (e.g., being scared of attending a particular event). In one condition, the bunny's tummy ache most frequently occurred in conjunction with a physiological event (i.e., a cause consistent with children's prior knowledge), whereas in the another condition, it co-occurred most frequently with a non-physiological event (i.e., a psychosomatic cause inconsistent with children's prior knowledge). Following the principles of Bayesian inference, children more readily accepted an event as the cause of a physiological state when evidence for that cause was consistent with their prior belief (i.e., when the cause was physiological) than when it was not. From 4-years old, children also *revised* their beliefs (i.e., also attributed the bunny's tummy ache to a psychosomatic cause) after seeing evidence conflicting with their initial belief, which supports the view that children learn by adapting existing theories and conceptual structures when encountering new evidence.

Furthermore, cross-cultural and cross-species research has shown that preliterate and prenumerate Mayan adults are just as capable of reasoning with probabilities as Mayan school children and Western controls (Fontanari et al., 2014), and that even great apes are capable of inferring the composition of a population from a randomly drawn sample (Hanus & Call, 2014; Rakoczy et al., 2014; Tecwyn et al., 2017). Albeit not being a uniquely human skill, the capacity to draw powerful inductive inferences from observable patterns of probabilistic evidence therefore appears to be a fundamental building block of (human) learning, which has strong evolutionary roots and develops early in life independently of formal training or education.

## 1.3. To explore or not to explore: Children's early sensitivity to uncertainty

Apart from relying on evidence the environment presents them, children also learn by actively exploring and intervening on their surroundings, thus producing their own data to inform their inductive inferences. The question of what motivates children, or learners more generally, to engage in exploration has recently become a major topic of interest in fields as diverse as ethology, artificial intelligence, robotics and neuroscience. The most influential current accounts of learning suggest that exploration is motivated by a learner's *epistemic curiosity*, a motivational state driving learners to engage in exploration, solely with the aim of gaining new information (as opposed to obtaining material rewards or accomplishing external goals). According to these accounts, the primary goal of exploratory behaviours is to reduce the learner's uncertainty about the world and to improve their prediction of future events (e.g., Gottlieb et al., 2013; Kidd & Hayden, 2015; Oudeyer, 2018). Relatedly, information-gap theory (Loewenstein, 1994) postulates that curiosity arises when an individual becomes aware of a gap in their knowledge, that is, when they are uncertain about something. Awareness of this knowledge gap induces a desire to reduce this gap – which is resolved by seeking the missing information.

Supporting these ideas, a growing body of literature demonstrates that children preferentially engage in exploration when there is more information to be gained, with this preference being present from infancy throughout childhood. Looking-time studies for instance have show that infants preferentially look at stimuli that are novel (Fagan, 1970; Fantz, 1964; Saayman et al., 1964), perceptually salient (Kaldy & Blaser, 2013), neither too easy nor too complex to process (Kidd et al., 2012, 2014), or that violate their intuitive expectations about the world (for a comprehensive review, see Spelke & Kinzler, 2007). In one study (Spelke et al., 1992), 2.5-months-olds were shown a ball rolling down a ramp and disappearing behind a screen. Partially visible behind the screen was a solid wall blocking the ball's path. When the screen was removed, infants either saw a possible outcome (i.e., the ball had been stopped by the wall) or an impossible outcome (i.e., the ball was at the lowest part of the ramp as if it had passed right through the wall) and looked significantly longer at the impossible outcome - presumably because they had intuitively expected the wall to stop the ball. Similar results that infants look longer at unexpected than expected events have been demonstrated in a range of knowledge domains, including physics (Baillargeon, 1987; Baillargeon et al., 1985; Kim & Spelke, 1992, 1999), numerosity (McCrink & Wynn, 2004; Wynn, 1992), agents and their goals, dispositions, and emotions (Kuhlmeier et al., 2003; Liu et al., 2017; Luo, 2011; Onishi & Baillargeon, 2005; Skerry & Spelke, 2014; Woodward, 1998; Wu & Gweon, 2019; Wu & Schulz, 2018), and social interactions (Hamlin et al., 2010; Powell & Spelke, 2013). While these findings have generally been taken to support the existence of intuitive theories or *core knowledge* in young infants (e.g., an intuitive understanding of objects solidity), they also suggest that surprise, or unexpectedness, may serve as a signal for potential learning opportunities, which young infants exploit by guiding their attention towards the unexpected event.

What's more, older infants not only look longer at surprising events, but also actively explore objects longer after being presented with evidence that violates their expectations (e.g., Sim & Xu, 2017a; Stahl & Feigenson, 2015): Stahl and Feigenson (2015) for instance showed 11- and 12-months-olds the solidity event described in Spelke et al. (1992) and then gave them the opportunity to manually explore the toy seen during the event or a novel distractor toy. Infants who had seen the expectation-consistent event (e.g., a ball stopped by a solid wall) preferentially explored the distractor toy (consistent with a general preference for novelty). However, infants who had witnessed the expectation-inconsistent event (i.e., a ball ostensibly moving through a solid wall) preferentially explored the familiar toy that had just violated their expectations about object solidity. Moreover, their exploration was clearly designed to test the property the toy had just violated, as infants tended to bang toys that had violated the principle of object solidity, but tended to drop toys that had violated gravity. In a similar study, 13-month-olds observed how an experimenter “randomly” drew samples of balls from two boxes containing differently coloured balls. The sample from the first box only contained balls of the same colour (an unexpected outcome given the mixed composition of the population), whereas the sample from the second box contained several differently coloured balls (an expected outcome given the population). When given the opportunity to explore both boxes, infants spent significantly more time touching and reaching into the box that had produced the unexpected sample (Sim & Xu, 2017a).

In addition, increased exploration following surprising events does not seem to merely reflect generally increased arousal but an actual drive to understand the unpredicted event: Perez and Feigenson (2022) showed 11-months-old infants the same solidity events as in Spelke et al. (1992) and Stahl and Feigenson (2015), but, before allowing them to explore the familiar and distractor toys, provided infants with an explanation for why the toy had passed through the wall (e.g., the back of the wall contained an opening large enough for the ball to move through). This effectively eliminated infants’ preferential exploration of the expectancy-violating toy, suggesting that they now treated the previously surprising event like an expected one and no longer sought an explanation for what they had observed.

Children’s preference for exploring under conditions of epistemic uncertainty is also present during the preschool years. For example, preschoolers selectively explore when evidence violates their prior beliefs (Bonawitz et al., 2012; van Schijndel et al., 2015), or when evidence is confounded (i.e., fails to disambiguate the cause of an effect; Cook et al., 2011; Gweon & Schulz, 2008; Lapidow & Walker, 2020; L. Schulz, Gopnik et al., 2007). Recent work also demonstrates that 4- and 5-year-olds are more likely to explore when alternative hypotheses are harder to distinguish, such as when the data produced by their own interventions is less discriminable (Siegel et al., 2021), which suggests that children’s exploration is not only sensitive to uncertainty arising from environmental patterns of evidence, but also with respect to the quality of their own interventions.

Around the age of 5, children also begin to show an *explicit* metacognitive understanding of epistemic uncertainty. In a recent study for example, 5- and 6-year-olds were more likely to state that they need more information to disambiguate two potential causal structures after being presented confounded (as opposed to unconfounded) evidence, and even began to accurately describe the reason for their insufficient knowledge (i.e., evidence failing to disambiguate a causal effect; Köksal et al., 2021), suggesting that they had a clear understanding of the uninformativeness of confounded evidence. However, it remains unclear to what extent such an explicit knowledge is related to children's exploration, with recent evidence indicating that a better metacognitive understanding may not necessarily be related to children's ability to perform informative interventions (Moeller et al., 2022).

## 1.4. Deciding how to explore: The efficiency and adaptiveness of children's active learning strategies

The reviewed body of work powerfully demonstrates that children play an active part in constructing their own knowledge. On the one hand, they are highly sensitive to information the environment presents them and draw powerful inferences about the world from probabilistic patterns and regularities they observe. On the other hand, children also actively engage in exploration in order to gain new data to support these inferences. Furthermore, they are highly skilled at detecting learning opportunities and preferentially explore in situations of epistemic uncertainty (e.g., when encountering conflicting, surprising or ambiguous evidence). However, deciding *whether or not to explore* is only the first step in successful active learning. Once the decision to explore has been made, any active learner also needs to decide *how to explore*. Considering that in any given situation, there is a potentially infinite number of actions to choose from, each with many possible outcomes, costs and rewards, deciding which action to perform is not a trivial problem. What's more, one action that is efficient at producing the sought information in one context may not necessarily be so in a different context. To illustrate, imagine your partner was outside in the garden and asked you to fetch him his phone. Having lived together for several years, you know that he always plugs his phone into its charger in the hallway the moment he comes home. Being fairly certain about where to find the phone, the most efficient action would be to head straight to the hallway to look for it. Now, imagine your partner had asked you to fetch him his wallet instead. Unfortunately, unlike his phone, he does not keep his wallet in a fixed place and you have no idea where he might have left it (e.g., in his work bag, in his sports bag, in the pocket of yesterday's trousers, on his desk, on the kitchen counter, or the back of the living room couch?). With there being so many potential locations, the more efficient search approach in this situation would be to ask your partner where he has left his wallet before attempting to fetch it, in order to reduce the number of locations you need to check before finding it.

As this example demonstrates, the usefulness, or informativeness, of any given strategy (e.g., attempting to find an object right away vs. first narrowing down the number of potential search locations) cannot not be defined in absolute terms. Instead, it depends on the particular characteristics of the task at hand (e.g., did your partner ask you to fetch the phone or the wallet?), as well as on previous knowledge and expectations of the searcher (e.g., knowing that your partner always plugs his phone into the charger; Todd & Gigerenzer, 2012). Moreover, even a randomly selected action is likely to lead to some learning at least (i.e., even checking any random location would tell you where the wallet is not). For learning to occur efficiently rather than accidentally therefore, children need to possess a mechanism to evaluate and weigh these different options in order to identify and select the action that is most likely to yield the desired information *given the current learning context* - an ability which has been referred to as *ecological learning* (Ruggeri et al., 2015; Ruggeri et al., 2017).

### 1.4.1. Learning from one's own interventions: Hypothesis-testing during free play

Despite children's strong abilities to detect learning opportunities based on degrees of uncertainty in the environment, evidence about the efficiency of children's active learning strategies is still mixed. Studies investigating preschoolers' exploratory strategies in causal learning tasks suggest that they are able to spontaneously design informative interventions both in experimental settings and during free play. In fact, there is evidence that children's exploration can be described as a simple form of hypothesis-testing, allowing them to confirm or disconfirm prior beliefs or to understand causal systems (e.g., Kushnir & Gopnik, 2005; Legare, 2012; Muentener & Schulz, 2014; Sobel & Sommerville, 2010; van Schijndel et al., 2015). For instance, in one study, 2-year-olds spontaneously pushed a button to test if it was connected to a light in the absence of a plausible cause for the light turning on (i.e., the light turned on randomly), but not when a plausible cause had previously been presented (e.g., an experimenter turned it on), suggesting that pushing the button was an active test of the hypothesis that the button caused the light to activate (Muentener & Schulz, 2014). Another study tested children (4- to 6-year-olds) who held the belief that only the size of an object determines the size of its shadow (rather than its size and distance from a light source) and presented them with evidence that either confirmed (i.e., a small puppet creating a smaller shadow than a large puppet) or conflicted this belief (i.e., a small puppet creating a larger shadow than a large puppet). In a subsequent free play period, children in the conflicting condition were more likely to spontaneously perform an unconfounded experiment, varying only one variable at a time to determine its effect, than in the confirming condition (van Schijndel et al., 2015). This indicates that preschoolers possess an intuitive understanding of basic principles of effective inquiry, such as the control-of-variables strategy (CVS), which consists of changing variables one by one to disambiguate their causal effects.

More recently, Lapidow and Walker (2020), presented 4- to 6-year-olds with a causal system that could operate according to two possible causes, and asked them to disambiguate the causal structure by performing one of two possible interventions. In particular, children were shown a gear toy composed of two interlocking gears, which was activated by pressing a button (causing the gears to spin). Children were then told that the toy was broken (i.e., the gears failed to spin), and that this could be either because *both gears* were broken and failed to spin (i.e., "multiple cause" structure), or because *one particular gear* was broken and prevented the second gear from spinning (i.e., "single cause structure"). Children were then given the choice of removing one of the two gears to test if the remaining gear would activate on its own. Crucially, there was precisely one informative action allowing children to determine which of the two causal structures was true: removing the gear that was broken in *both structures* to observe if the remaining gear would activate on its own. The majority of children successfully removed the gear that could disambiguate between the two causal structures, suggesting that they were sensitive to the informative potential of their own causal interventions.

Importantly, there is also some evidence that children's interventions actually result in learning (Bonawitz et al., 2011; Lapidow & Walker, 2020; McCormack et al., 2016; L. Schulz, Gopnik et al., 2007; Shneidman et al., 2016; Sobel & Sommerville, 2010). In the study by Lapidow and Walker (2020) for instance, almost all tested children were able to accurately infer whether the single or the multiple cause structure was true after performing the informative intervention of removing the gear broken in both structures. In another causal learning study with 2- and 3-year-olds (Sim & Xu, 2017b), children were asked to infer whether different blocks activated a machine based on a shape rule (i.e., blocks matching the machine's shape activated the machine), or a colour-rule (i.e., blocks matching the machine's colour activated the machine). Children were as successful at activating a familiar and a novel machine after a free-play period as after a didactic demonstration by an experimenter, suggesting that children were able to generate the necessary evidence on their own, could use that evidence to determine how to activate the familiar machine, and could generalise their knowledge to a novel problem. What's more, in a subsequent study using the same paradigm but with a more complex causal structure (i.e., two blocks, not one, matching the machine's colour, or the machine's shape had to be placed on the machine to activate it), 3- and 4-year-olds actually performed better following a period of free play than following didactic instruction (Sim et al., 2017, also see Yuniarto et al., 2020).

On the other hand, there is also some evidence that children's learning from their own interventions is limited. In the study by van Schijndel et al. (2015) for instance, children in the conflicting condition displayed a strong preference for varying the size of the puppet, rather than its distance from the light source. This suggests that they preferentially designed experiments to confirm, rather than to falsify their theories (a tendency also commonly observed in studies investigating children's scientific reasoning; Klahr et al., 1993; Schauble et al., 1991; Wason, 1960), and that observing conflicting evidence was not sufficient to make children consider and test an alternative hypothesis, namely that distance also affects shadow size. What's more, the learning observed following children's free-play period was almost insignificant, with only 10% of children progressing towards the more advanced understanding that both size *and* distance determine the size of an object's shadow. It is unclear why children adhered to their initial hypothesis. Bayesian models of learning would suggest that the counter-evidence they observed was not strong enough to override their initial belief. Alternatively, children may have lacked the necessary cognitive flexibility to consider an alternative hypothesis, or may have drawn the wrong conclusions from their experiments. More generally, these findings highlight one crucial limitation of children's ability to learn from free play, namely that their interventions can only lead to learning when the appropriate hypotheses are considered (but see Lapidow and Walker, 2022 for an alternative perspective that even experiments designed to confirm theories may lead to learning).

In addition, educational findings on children's scientific reasoning skills demonstrate that even primary-school aged children struggle to employ informative inquiry strategies, such as the control-of-variables strategy, when exploring more complex scientific phenomena, and that use of CVS does not develop spontaneously but requires intensive training (for a recent review, see Shtulman & Walker, 2020). Likewise, information-board studies investigating children's predecisional information search typically find that children's information search is largely inefficient, unsystematic and exhaustive, with particularly great improvements observed between the ages of 5 and 10 (Betsch et al., 2018; Davidson, 1991a, 1991b, 1996; Gregan-Paxton & John, 1995, 1997; Howse et al., 2003; Lang & Betsch, 2018; Lindow & Betsch, 2019), although a recent study using a simplified predecisional information-search paradigm has revealed more efficient search strategies in children as young as 4 (Lindow, 2021).

### 1.4.2. Children's question asking

Another line of research investigating the efficiency of children's exploration has focused on children's questions. Questions are one of the most powerful tools for children to learn about the world as they allow them to request exactly the piece of information they are missing, at the exact moment they need it most in order to resolve a state of uncertainty (Chouinard et al., 2007). What's more, by asking questions, children can learn about concepts they would not be able to learn about through mere observation or self-directed exploration, such as cultural, religious, or scientific concepts (e.g., Harris, 2000; Harris & Koenig, 2006). While during the first year of life, children primarily ask about facts or object labels, the questions of older children target a large variety of topics, including their own behaviour, the behaviour of others, cultural and social conventions, natural phenomena, and physical mechanisms (e.g., Callanan & Oakes, 1992; Chouinard et al., 2007; Hickling & Wellman, 2001; Kelemen et al., 2005). Children's questions also become increasingly complex between the age of 2 and 5 (for current reviews, see Jones, Swaboda et al., 2020; Ronfard et al., 2018), and there is strong evidence that their questions are purposeful, intended to fill knowledge gaps, and to test and extend their developing understanding of the world (e.g., Carey, 1985; Chouinard et al., 2007; Frazier et al., 2009; Harris, 2012). For example, children begin to seek causal explanations around the age of 2 (Callanan & Oakes, 1992; Chouinard et al., 2007; Hickling & Wellman, 2001), and have reasonable expectations about what constitutes a satisfying answer to their questions by the age of 3, as indicated by their tendency to ask more questions or to provide their own answers when receiving a non-informative reply (e.g., Frazier et al., 2009, 2016; Kurkul & Corriveau, 2018). Previous research has also shown that 4- and 5-year-olds ask domain-appropriate questions (e.g., asking about the functions of artefacts, and about eating behaviours of animals; Greif et al., 2006), and direct their questions to more reliable informants (for reviews, see Mills, 2013; Sobel & Kushnir, 2013). What's more, from around 4 years of age, preschoolers are able to generate questions that are mostly informative, as opposed to uninformative, irrelevant or redundant, and by the age of 5, they reliably use the answers they receive to solve problems

(Legare et al., 2013).

Despite these findings that children's questions are meaningful and support learning, there is also contrasting evidence suggesting that young children have great difficulty asking informative questions. Kemler et al. (2004) for example presented 3- and 4-year-olds with novel artefacts and gave them an opportunity to ask about them. They found that children's questions were often ambiguous, that is, did not clearly communicate whether children wanted to learn about the artefacts' functions or names (e.g., "What is it?"). What's more, children were more likely to ask follow-up questions after receiving the artefact's name, suggesting that they actually had a preference for learning about the function of these artefacts, but failed to communicate this preference in their questions.

Preschoolers' difficulties with formulating those questions that allow them to obtain the information they seek has also been documented in studies using the 20-questions game. The goal of this game is to identify a target object or cause (e.g., "What objects can be found on Planet Apres?" or "Why was the man late for work today?") within a given set of candidate hypotheses by asking as few yes–no questions as possible (e.g., Herwig, 1982; Legare et al., 2013; Meder et al., 2019; Mosher & Hornsby, 1966; Ruggeri & Feufel, 2015; Ruggeri & Lombrozo, 2014; Ruggeri et al., 2015, 2016; Ruggeri et al., 2017). The game is solved by sequentially narrowing down the given set of hypotheses or candidate solutions ("hypothesis space"), until only one hypothesis, the correct solution, remains. Although the 20-questions game may appear rather constrained and artificial at first glance, it is a classic example of sequential, binary information search, a problem that is actually a very general one encountered throughout the lifespan, such as in real-world decision making, categorisation, causal inference and medical diagnosis (e.g., Berretty et al., 1997; Berretty et al., 1999; Green & Mehr, 1997; Hamilton, 2003; Martignon et al., 2008). The 20-questions game therefore provides a good compromise between experimental tractability and real-world generalisability.

There are two possible approaches to solving a 20-questions game (Mosher & Hornsby, 1966): *hypothesis-scanning* questions are questions that test single hypotheses within the set of candidate solutions directly (e.g., "Is this daisy on Planet Apres?" or "Was the man late because he woke up late?"). In contrast, *constraint-seeking* questions target a category of objects or features shared by several hypotheses (e.g., "Are flowers found on Planet Apres?" or "Was he late because he forgot something at home?"). Constraint-seeking questions are usually considered more efficient than hypothesis-scanning ones because they allow the question-asker to rule out multiple hypotheses (objects, categories of objects or reasons) at once, thus reducing the number of questions needed to identify the solution (Mosher & Hornsby, 1966). Studies employing the 20-questions game typically find that children exclusively ask hypothesis-scanning questions until about age 7, and that the proportion of constraint-seeking questions increases over the primary school years, until constraint-seeking becomes the dominant approach in adulthood (Herwig, 1982; Mosher & Hornsby, 1966; Ruggeri & Feufel, 2015; Ruggeri et al., 2015, 2016). For example, Herwig (1982) found that in a 20-questions game, only 5% of questions asked by preschoolers, 10.5% of those asked by first graders, 17% of those asked by second graders, and about 80% of those asked by fifth graders were constraint-seeking questions. In addition, although even 4-year-olds are able to spontaneously generate constraint-seeking questions, their constraint-seeking questions are often not the most informative ones available (Legare et al., 2013; Mills et al., 2010; Ruggeri et al., 2021).

Apart from a qualitative distinction between (more informative) constraint-seeking questions, and (less informative) hypothesis-scanning questions, researchers have recently also employed more quantitative measures of informativeness to determine the efficiency of different questions, such as expected stepwise *information gain* (see Chapter 2.3). These studies have generally confirmed that the efficiency of children's question-asking strategies steadily increases over the childhood years (Kachergis et al., 2017; Meder et al., 2019; Nelson et al., 2010; Ruggeri & Feufel, 2015; Ruggeri & Lombrozo, 2014; Ruggeri et al., 2016), and that preschoolers (4- to 6-year-olds) in particular have great difficulties generating informative questions from scratch (Ruggeri et al., 2021).

#### 1.4.2.1 Why do children fail to ask informative questions?

There are likely several potential factors contributing to the observed developmental changes in the efficiency of children's question-asking strategies, including verbal skills, categorization skills, executive functions, metacognition, probabilistic reasoning, attention, motivation, education, and socioeconomic status (see Jones, Swaboda et al., 2020, for a review). One dominant explanation for young children's tendency to ask hypothesis-scanning questions is that constraint-seeking questions are more cognitively demanding than hypothesis-scanning ones because they require the ability to identify and mentally represent the various dimensions into which the items in the hypothesis-space could be categorised, for example based on shared features. Supporting this idea, Legare and colleagues (2013) found that the ability to identify and flexibly categorize objects based on alternative features (e.g., colour vs. pattern) predicted how well 4- to 6-year-olds generated informative questions. What's more, there is also evidence that priming 7- to 10-year-old children to represent the hypothesis space in terms of its abstract categories rather than in terms of individual hypotheses can facilitate their ability to ask constraint-seeking questions: Ruggeri and Feufel (2015) presented children and adults with word labels for 20 different objects (e.g., "dog" or "sheep"), describing them either at the more abstract basic level (e.g., "dog") or at the object-based subordinate level (e.g., "dalmatian"). Participants were more likely to ask constraint-seeking questions in the basic-level condition compared to the subordinate condition, suggesting that the more abstract labels helped children to move away from object-based reasoning, thereby allowing them to go beyond the hypothesis-scanning approach, when generating their questions. Furthermore, this study found that being able to actively produce more abstract features (e.g., "a dog is a mammal" or "a dog has four legs") was positively related to children's question-asking performance, which supports the idea that improvements in categorisation skills are one driver behind the observed developmental changes in question-asking efficiency. Finally, it has been shown that preschoolers are more likely to ask constraint-seeking questions when they do not need to generate questions from scratch but can select among several given questions (Ruggeri et al., 2017), potentially because this eliminates the need to identify and generate those categorical dimensions that can be used to structure and partition the hypothesis space.

A second possible explanation for young children's difficulties with formulating constraint-seeking questions is that these questions may be more linguistically challenging than hypothesis-scanning ones. In particular, asking constraint-seeking questions requires children not only to be familiar with and produce the labels of the different categories and their distinguishing features, but also to flexibly produce various question structures (e.g., "Is it ...?", "Does it have...?", "Can it be used for...?", "Is it made of...?"). In contrast, a hypothesis-scanning question can be asked using a single grammatical question form (e.g., "Is it ...?"), only requires children to know the labels of the candidate hypotheses, and can even be made by simply pointing at the available solutions. Supporting this idea, there is evidence that efficient information search emerges much earlier in tasks that do not involve verbal question asking (e.g., Bonawitz et al., 2012; Cook et al., 2011; Kushnir & Gopnik, 2005; E. Schulz et al., 2019; L. Schulz, Gopnik et al., 2007). However, the link between verbal skills and performance on the 20-questions game has never been investigated directly. This is one aim of the present dissertation (see below).

### 1.4.3. Adaptiveness and ecological learning

As documented by the studies reviewed in the previous section, there are large developmental improvements in children's ability to select those questions or exploratory strategies that are most likely to provide them with the sought information. However, as mentioned above, the efficiency of different queries cannot be defined in absolute terms, but actually depends on the characteristics of the particular learning context and on the previous knowledge and expectations of the searcher (Todd & Gigerenzer, 2012). For active learning to be truly efficient, learning strategies must therefore also be adaptive, tailored to the *current learning environment*. The first to demonstrate that children are such adaptive and *ecological learners*, capable of dynamically selecting those strategies that provide the largest information gain, were Ruggeri et al. (2015). They tasked 7- and 10-year-old children and adults to find out why a man was late for work and provided them with 10 candidate solutions. In one condition, all hypotheses were presented as equally likely (uniform condition). In the other condition, a few hypotheses were introduced as much more likely to be the correct solution than the others (mixed condition). As described above, the more efficient approach in this scenario consists of testing the single most likely candidate solution directly by asking a hypothesis-scanning question, and not of asking a constraint-seeking question. Although, replicating earlier findings, the overall rate of constraint-seeking questions steadily increased with age, all age groups, including 7-year-olds, were more likely to ask hypothesis-scanning questions in the mixed condition than in the uniform condition, suggesting that they had identified the hypothesis-scanning approach as more efficient in the mixed condition and adapted their questions accordingly.

In another 20-questions study (Ruggeri et al., 2017), 5-year-olds had to find out why the monster Toma had been late for school. Participants were initially presented a storybook describing why Toma had been late for school over several days. In the uniform condition, Toma had been late equally often for six different reasons. In the skewed condition, Toma had been late multiple times for one particular reason (i.e., single most likely hypothesis), and three times each for one different reason. Children then learned that Toma had been late yet again and that two of his monster friends wanted to find out why. In the uniform condition, one of the monster friends asked a constraint-seeking question, targeting half of the available candidate solutions (i.e, the optimal constraint-seeking question), whereas the other friend asked a hypothesis-scanning question targeting only one of the available candidates. In contrast, in the skewed condition, the first monster friend asked a constraint-seeking question targeting 3 out of 8 hypotheses, whereas the other monster friend asked the hypothesis-scanning question that targeted the single most likely candidate solution. Hence, in the two conditions, the two question types differed in their relative informativeness, with the hypothesis-scanning question more informative than the constraint-seeking question in the skewed condition, but less informative in the uniform condition. When asked which of Toma's friends would find out first why Toma had been late for school yet again, children in both conditions selected the monster asking the more informative question irrespective of question type: 70% of the children selected the constraint-seeking friend in the uniform condition, and 73% of the children selected the hypothesis-scanning friend in the skewed condition. What's more, children did so dynamically, such that, when playing the game over several rounds, they reassessed the informativeness of the different questions and adapted their choices accordingly. These results therefore demonstrate that children as young as 5 are ecological learners, able to dynamically select the most informative questions in the current context.

Recently, the adaptiveness of children's exploratory actions has also been assessed in the domain of causal learning and children's use of the control-of-variables strategy. Building on findings that the informativeness of CVS depends on the degree of casual sparsity, that is, the number of variables likely to affect the outcome (Coenen et al., 2019), Jones and colleagues (2020) presented 7- to 13-year-old children with a fusebox with 6 switches. In one condition, children were told that one switch turned on the light (sparse condition), whereas in the other condition, they were told that all switches turned on the light when flipped individually, but that one was broken (non-sparse condition). In both conditions, they were asked to identify the unique switch (i.e., the single working one in the sparse condition, and the single broken one in the non-sparse condition). Mirroring findings from the 20-questions game that younger children preferentially ask hypothesis-scanning questions, this study found that children overall preferred to test switches individually, but that this tendency decreased with age. However, and consistent with the finding that testing multiple candidate variables at once (as opposed to testing variables one by one) becomes more efficient as causal sparsity increases (i.e., as fewer variables affect an outcome, Coenen et al., 2019), children as young as 7 preferentially tested multiple switches at once in the sparse condition as compared to the non-sparse condition, indicating that their choice of strategy was sensitive to the causal sparsity of the presented problem.

To date, little research has assessed whether children younger than 5 are ecological learners. Considering that a strong sensitivity towards environmental probabilities seems to be present from infancy (see Section 1.2.), it seems likely that ecological learning may already be present in much younger children. However, due to their complexity (causal sparsity task) or reliance on linguistic skills (20-questions game), the paradigms previously used to study ecological learning do not seem appropriate for preschoolers or toddlers. It is the second aim of this dissertation to investigate the ecological learning capacities of children below the age of 5 using more age-appropriate paradigms.

## 1.5. The present dissertation

The work reviewed in this chapter leaves two somewhat contradicting conclusions: on the one hand, children are effective active learners, highly attuned to learning opportunities and capable of designing informative interventions from a very young age. On the other hand, efficient information search only seems to emerge during the primary school years. How can these contrasting lines of evidence be reconciled? A first potential explanation may lie in the differences in paradigms used to investigate children's exploration. Evidence for children's early active learning competence comes primarily from causal learning studies (e.g., Cook et al., 2011; Kushnir & Gopnik, 2005; Lapidow & Walker, 2020; Legare, 2012; Muentener & Schulz, 2014; Perez & Feigenson, 2022; Sim & Xu, 2017a; Sobel & Sommerville, 2010; Stahl & Feigenson, 2015). These studies typically consider the *effectiveness* of children's exploratory actions, that is whether their actions result in the intended outcome or not, and look at whether children successfully perform *one particular effective action* or not (e.g., dropping or banging a toy, varying the size or distance of a puppet from a light source, removing the informative rather than the uninformative gear). However, they do not take into account that exploratory actions may have varying levels of informativeness, with some being *more*, and some being *less efficient* at producing the intended outcome.

In contrast, evidence for late information-search competence has primarily been derived from the 20-questions game, although late emergence of efficient inquiry has also been demonstrated in the domain of scientific reasoning (Shtulman & Walker, 2020) and predecisional information search (Betsch et al., 2018; Davidson, 1991a, 1991b, 1996; Gregan-Paxton & John, 1995, 1997; Howse et al., 2003; Lang & Betsch, 2018; Lindow & Betsch, 2019). Crucially, in the 20-questions game, there are two information-search strategies available, which are both effective (i.e., lead to the solution) but differ in their *relative expected efficiency* (i.e., their expected stepwise information gain), and conclusions about children's search inefficiency have been derived from the fact that they tend to use the *less efficient*, but nonetheless effective, approach. The 20-questions paradigm therefore assesses children's exploration in a more fine-grained way by investigating whether they chose *the more informative of two available strategies* (as opposed to implementing one single informative action or not). On the other hand, as described above, the 20-questions game relies heavily on children's verbal skills, and therefore seems ill-suited for studying the efficiency and adaptiveness of the exploratory strategies of preschool-aged children, whose linguistic skills are still developing.

Given these considerations, the present dissertation has two primary aims: a) to extend the investigation of children's ecological learning capacities to children younger than previously tested, and b) to investigate the efficiency of preschoolers' information search in tasks that do not require question generation. More specifically, it addresses the following questions:

1. Are children younger than 5 years adaptive, ecological learners? That is, do they flexibly select and implement those exploratory strategies that promise the largest information gain given the current learning context?
2. Do preschoolers search for information more efficiently in non-verbal tasks, such as in a spatial search task?
3. What are potential reasons for the inefficiency of preschoolers' questions in the 20-questions game?
4. Do children (and adults) disentangle *expected efficiency* of exploratory actions from their *observed effectiveness*?

To investigate these questions, this dissertation introduces a range of novel non-verbal paradigms designed following the formal and computational structure of the 20-questions game, which has been the primary tool for studying the efficiency and adaptiveness of children's information search, but entirely removing the question-asking component from it. By extending this formal framework, which allows quantification and comparison of the efficiency of children's exploratory actions using expected stepwise information gain (see Section 2.3), to preschoolers and toddlers, this dissertation attempts to bridge the gap between work on children's causal learning, suggesting early active learning competence, and findings from the 20-questions game, suggesting that information-search competence only emerges later in childhood. Ultimately, in providing evidence for, or against, the idea that children are efficient and adaptive active learners, this dissertation also intends to bring new insights into the debate of whether play and self-directed exploration are suitable means for educating young children, or whether they should better be taught through direct instruction to maximise learning outcomes.

# 2. Methodology

## 2.1. Samples

As the main focus of this dissertation lies on children below primary school age, participants' age in the reported studies ranged from 2 to 5, although Study 2 additionally included 6- and 7-year-olds, and Study 5 also included children up to the age of 10. To ensure more socio-economically diverse samples, children were recruited and tested across several different museums, as well as one preschool, in Berlin, Germany. All participants were fluent in German, and had no learning disabilities. Note that the sample of Study 1 was partially collected in the US. For these children, fluency in English was a prerequisite for participation. Recruitment and testing of this subset also took place at local museums. Prior to participation, written informed consent of legal guardians was obtained for all participants. Children were asked for verbal consent and received small gifts (e.g., stickers) as a reward for taking part in the studies. All studies were approved by the Ethics Committee of the Max Planck Institute for Human Development, Berlin.

## 2.2. General methodological approach

In order to answer the research questions outlined above, this dissertation employed a behavioural approach combined with computational measures quantifying the informativeness of the different exploratory options available to children in the developed tasks. In particular, it introduced a range of novel non-verbal behavioural tasks, which were based on the formal and computational structure of the 20-questions game but removed the question-asking component from it, thereby allowing the investigation of the efficiency and adaptiveness of the information-search strategies of much younger children, and in domains other than question asking. Importantly, in keeping the formal and computational structure of the 20-questions game, and quantifying the efficiency of the available exploratory actions using expected stepwise information gain (see below), it was possible to make clear and empirically testable predictions as to which action or query children should chose in order to maximise the information gained through their interventions. An additional advantage of this approach is that performance of children tested in the non-verbal studies presented here can be more easily put into continuity with earlier findings from the original 20-questions game, and it is precisely through this continuity that this dissertation aims to bridge the gap between findings of early versus late emergent information-search competence documented in previous work.

This dissertation is built on two projects. The first paper (Swaboda et al., 2018), summarised in Chapter 2.4.1, focuses on the adaptiveness of children's exploratory strategies in a task where they had to find an egg shaker hidden in one of four small boxes contained within two larger boxes and could select among two exploratory actions (opening vs. shaking the boxes first) to retrieve the egg. This experiment varied how likely the egg was to be found in each of the four small boxes (skewed vs. uniform condition), so that the two actions were differently informative in the two conditions. The second paper (Swaboda et al., 2022), summarised in Chapter 2.4.2, explores the efficiency of children's information-search in a spatial search task, in which the task's hierarchically organised hypothesis-space was visually laid out as a two-dimensional maze, and children had to discover the path through that maze. Queries were made by removing cards covering the various nodes of the hierarchy (corresponding to the passages of the maze, which could either be open or closed). Project 2 additionally compared children's search strategies in this spatial search task to a hierarchical version of a traditional 20-questions game, where they had to identify a target monster (which could activate a music machine) from a set of eight candidate monsters. For copyright reasons, the published manuscripts are not provided in this dissertation, but they can be accessed following the links in their respective reference:

**Paper 1:** Swaboda, N., Ruggeri, A. & Gopnik, A. (2018). Preschoolers adapt their exploratory strategies to the information structure of the task. *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*, 1102–1107. https://cogsci.mindmodeling.org/2018/papers/0218/0218.pdf

**Paper 2:** Swaboda, N., Meder, B. & Ruggeri, A. (2022). Finding the (most efficient) way out of a maze is easier than asking (good) questions. *Developmental Psychology*, Advance online publication. https://doi.org/10.1037/dev0001386

In addition, I was involved as second and third author in four other projects (Chai et al., in preparation; Serko et al., in preparation; Török et al., under revision; Ziemann et al., in preparation), as well as in a book chapter exploring developmental changes in children's question asking (Jones, Swaboda et al., 2020). As these projects are highly relevant to the questions addressed in this dissertation and will be referenced in the Discussion, a short summary of each of them is provided in Chapter 2.5.

## 2.3. Computational framework

In both published projects, the efficiency, or informativeness, of children's queries was quantified using the computational measure of expected stepwise *information gain* (e.g., Eimas, 1970; Kachergis et al., 2017; Nelson et al., 2014; Ruggeri & Feufel, 2015; Ruggeri et al., 2015, 2016), which assesses how much each query reduces the uncertainty, measured via Shannon (1948) entropy, in the hypothesis space (see Chin et al., 2015; Meder et al., 2019; Nelson et al., 2010; Oaksford & Chater, 1994; Steyvers et al., 2003). Within this framework, the most efficient questions are those that maximise the expected reduction of entropy, allowing the learner to move from a state of uncertainty (e.g., "Which box contains the egg shaker?", "What is the correct path through the maze?", "Which one is the special monster?") closer to a state of certainty (e.g., "The egg shaker is in the yellow box.", "The path through the maze goes through here.", or "The special monster is round, yellow and has dots"). Expected stepwise information gain is calculated as the difference between the prior entropy, i.e. the uncertainty *before* making a query, minus the expected posterior entropy *after* making that query. The expected posterior entropy takes into account the entropy of all future possible states (e.g., the number of small boxes, possible paths through the maze, or candidate monsters), as well as their respective likelihoods (i.e., how likely it is to obtain a yes or no answer).

A great advantage of this quantitative approach over a purely qualitative one is that it allows researchers to compute the exact expected information gain of every conceivable (task-related) query and thus provides a much more fine-grained measure of the informativeness of different questions (e.g., see Jones, Swaboda et al., 2020). In particular, expected stepwise information gain can capture the varying informativeness of different constraint-seeking questions, as well as task-dependent variations in the relative informativeness of hypothesis-scanning and constraint-seeking questions. In a 20-questions scenario where all candidate hypotheses are equally likely (i.e., when there is a *uniform* distribution of likelihoods across all potential solutions), stepwise information gain is directly related to the split induced by a query, with those questions that induce the highest split (i.e., come closest to a 50:50 split) having the highest information gain (Meder et al., 2019; Navarro & Perfors, 2011; Nelson et al., 2014; Nelson et al., 2018). In Project 2 (Swaboda et al., 2022) for example, the initial hypothesis space in both games consisted of eight alternative hypotheses, corresponding to the eight monsters in the 20-questions game and the eight possible maze configurations in the maze-exploration game. Assuming a uniform prior (i.e., all hypotheses are equally likely), the prior entropy at each step is $\log n$, where $n$ is the number of remaining items.[1] The expected information gain ($IG$) of a question or query $Q$ is

$$\mathrm{IG}(Q) = \log n - \left[\frac{n_{yes}}{n}\log n_{yes} + \frac{n_{no}}{n}\log n_{no}\right] \tag{2.1}$$

where $n_{yes}$ indicates the number of hypotheses that would remain under consideration in case the answer were affirmative (i.e., if the target monster has the queried feature value, or if the queried passage is open), and $n_{no}$ indicates the number of hypotheses that would remain under consideration in case the answer were negative (i.e., if the target monster does not have the queried feature value, e.g. is not green, or if the queried passage is closed). The term in brackets is the expected uncertainty (entropy) after having asked the question. When all candidate hypotheses are equally likely, the maximum expected stepwise information gain value is 1, and is given by those queries that split the hypothesis space exactly in half, as such queries maximally reduce uncertainty. In contrast, expected stepwise information gain decreases as the split induced by a query becomes less even, with the minimum possible value being 0 (no information gathered and no reduction of uncertainty, e.g. confirmatory or irrelevant questions). In both tasks, first queries targeting the bottom level (i.e., one of the eight individual monsters or maze exits) induce a 1:7 split (equivalent to a hypothesis-scanning query, *IG* = 0.544), whereas first queries targeting the middle or highest level induce a 2:6 split (equivalent to a suboptimal constraint-seeking query, *IG* = 0.811) and a 4:4 split (equivalent to the optimal constraint-seeking query, *IG* = 1), respectively. An optimal searcher would select, at every step of the search process, the query with the highest expected stepwise information gain. Because of their hierarchical structure, the optimal search strategy in both tasks consists of a “top-down” approach, which first targets the highest, then the middle, and then the lowest level. Each of these queries induces a 50:50 split in the remaining hypothesis-space (4:4, 2:2, 1:1), and offers the maximum expected stepwise information gain of 1 at each step. Thus, this search approach guarantees finding the solution with the minimum possible number of three quer-

[1] The choice of base for the logarithm is arbitrary; we use base 2 in which case the unit is the bit.

ies.[2] The lower bound of search performance in both tasks is provided by a searcher who asks, at every step of the search process, the question with the lowest expected information gain (equivalent to asking only hypothesis-scanning questions). With a total of eight candidate hypotheses, this searcher would find the solution with a maximum of seven queries (i.e., ruling out hypotheses one by one until only the target remains), with the mean *IG* ranging from 0.54 (one query needed) to 0.75 (seven queries needed).

In contrast, in situations where one solution is proportionally much more likely than the others, such as in Project 1 (i.e., when the likelihood distribution is *skewed* towards one particular candidate solution; Swaboda et al., 2018), a hypothesis-scanning question targeting that single most likely hypothesis may actually be more informative than the best constraint-seeking question. Project 1 contrasts a uniform likelihood distribution (i.e., the egg could be in any of the four small boxes) with the most extreme case of a skewed distribution (i.e., the egg has a 100% chance of being in one particular small box). In this particular case, shaking and opening actually both have an *IG* of 0 as there is no prior uncertainty as to where the egg shaker is and uncertainty cannot be reduced further (for more probabilistic examples of skewed environments resulting in more graded variations in expected information gain, see Jones, Swaboda et al., 2020). In the skewed condition, shaking a large box first is therefore unnecessary to determine which box should be opened, and the most goal-directed action consists of opening the target box right away. On the other hand, applying equation 2.1, shaking one of the two large boxes first in the uniform condition induces a 50:50 split (*IG* = 1), while opening a small box right away induces a 1:3 split (*IG* = 0.81). Shaking one of the two larger boxes first is therefore the more informative action in the uniform condition.

Using such quantitative measures to study the informativeness of children's queries does not assume that children, or adults, actually mentally perform those complex computations when choosing an exploratory action or question. Instead, they likely rely on more superficial cues, such as the number of alternative solutions each query can rule out at once, or the evenness of the split induced by a query. Indeed, there is evidence that in the 20-questions game, the strategy of older children and adults can be characterised in terms of this split-half heuristic (Denney & Denney, 1973; Eimas, 1970; Nelson et al., 2014; Siegler, 1977; Thornton, 1982), suggesting that expected information gain provides a suitable computational model for the investigation of the efficiency of children's exploratory actions.

[2] Note that selecting queries with maximum information gain does not necessarily guarantee to identify the most efficient search path, i.e. sequence of queries (see Meder et al., 2019; Nelson et al., 2018 for details.)

## 2.4. Summary of associated published projects

### 2.4.1. Paper 1: Preschoolers adapt their exploratory strategies to the information structure of the task

The first publication was submitted to the conference organisers of the 40th Annual Meeting of the Cognitive Science Society in February 2018 and accepted for publication in the peer-reviewed conference proceedings in April 2018. As first author, my contribution to this publication consisted primarily of data analysis and writing the manuscript (70% of the work). Prof Ruggeri was responsible for the initial task design and supervised the statistical analyses and drafting of the paper (30%). Note that I was also involved in a follow-up on this project, which explored why 5-year-olds failed to demonstrate adaptiveness (see below), and has been published in the journal *Cognition* (Ruggeri et al., 2019). As third author of this publication, I was responsible for designing the follow-up, data collection, statistical analyses, and drafting the initial manuscript (60% of the work). Prof Ruggeri finalised the draft (20%) and Prof Gopnik (10%) provided valuable feedback to guide the development of the paper.

This project explored the adaptiveness of children's exploratory strategies using a novel non-verbal exploration paradigm. In particular, it investigated whether children as young as 3 years old are ecological learners, able to select those exploratory strategies that are most informative in a given task. Children ($N$ = 114; 3 to 5 years old) searched for an egg shaker hidden in one of four small boxes contained in two larger boxes. Prior to this game, children learnt that the egg was equally likely to be found in any of the four small boxes (uniform condition), or that it was most likely to be in one particular small box (skewed condition). Children then chose between two exploratory actions: they could either *open* a large and a small box right away, or *shake* one or two large boxes first to hear which one contained the egg. Importantly, these two actions were differently informative in the two conditions. Children in the uniform condition could not know in which box the egg was hidden. Shaking at least one large box before deciding which one to open was therefore necessary to avoid opening the wrong box. In contrast, children in the skewed condition should know where the egg was hidden and could open the correct large and small box right away. Results showed that 3- and 4-year-olds, but not 5-year-olds, tended to shake the boxes first in the uniform condition, but preferred to open a box right away in the skewed condition. This suggests that children correctly inferred how likely the egg was to be found across the four small boxes based on the frequency pattern they had observed, and, more importantly, that they tailored their exploratory actions to these different likelihood distributions. Hence, this project is the first to demonstrate that children as young as 3 are ecological learners, able to dynamically select those exploratory actions that provide the largest information gain in a given task.

A follow-up study, published in *Cognition*, further demonstrated that the failure of 5-year-olds to search adaptively may have been due conflicting verbal cues between the frequency training and test phase of the experiment. In particular, during frequency training, the experimenter said that she always *put* the ball in the same/different box. However, during the test phase, she said that she would *hide* the ball in one of the four small boxes. This may have led the older children to doubt the relevance of the frequency training for the test phase, who may instead have believed that the hiding game in the test phase had different rules from the one they had played earlier during training. As the results of this follow-up showed, 5-year-olds searched as adaptively as the younger children in the original study once these inconsistent cues were removed (Ruggeri et al., 2019).

### 2.4.2. Paper 2: Finding the (most efficient) way out of a maze is easier than asking (good) questions.

The second publication was submitted to *Developmental Psychology* in July 2020 and was accepted for publication in March 2022. As first author, I was responsible for 75% of the work involved in this publication, including development of the behavioural paradigm, data collection, statistical analysis, writing the manuscript, responding to reviews and implementing revisions. Dr. Björn Meder (10% of the work) and Prof Azzurra Ruggeri (15%) also contributed to the initial task design and the design of the requested follow-ups, and took on supervisory roles, providing feedback on statistical analyses and manuscript drafts.

This project explored reasons for the documented information-search inefficiency of younger children in the 20-questions game (e.g., Herwig, 1982; Mosher & Hornsby, 1966; Ruggeri & Feufel, 2015; Ruggeri et al., 2015, 2016). In particular, it investigated whether 4- to 7-year-old children's information-search strategies can be scaffolded and supported by a) highlighting the categorical relationships between hypotheses by representing them *spatially*, and b) reducing verbal demands associated with question generation. Children played a spatial search task, the *maze-exploration game*, in which they had to discover the path through a maze by uncovering its concealed passages to find out which passages were open and which ones were closed, and a traditional version of the 20-questions game, where they had to identify a target monster from a set of eight monsters by asking yes-no questions. The maze-exploration game was designed to be structurally and computationally analogous to the 20-questions game, thus allowing the comparison of children's search performance across both tasks. Just as in the 20-questions game, where the monsters varied across three hierarchically and symmetrically structured features, giving rise to eight unique candidate monsters, the maze was symmetrically organised onto three hierarchical levels, with two passages on the higher, four passages on the middle, and eight passages on the lower level. The results of Study 1 showed that children ($N$ = 124, 4- to 7-year-olds) employed more informative search strategies in the maze-exploration game than in the 20-questions game: they achieved higher mean information gain, and were more likely to make the most informative first query and to employ the overall optimal strategy. Indeed, the maze was so effective that it boosted younger children's search performance up to the level achieved by 7-year-olds. This result indicates that, under the right conditions, the ability to conduct informative queries may emerge at an earlier age than previously assumed.

A series of follow-up studies found that the search improvements observed in the maze-exploration game were primarily due to the fact that children could make queries non-verbally (i.e., by removing passages instead of asking questions): Study 2a introduced a feature-selection version of the 20-questions game, where children were presented with three *feature cards*, each depicting all variants of one feature, and could use these cards to select (i.e., point at) the feature and variant they wished to query instead of having to verbalise a question. Children ($N$ = 135, 4- to 6-year-olds), searched as efficiently in the feature-selection version of 20-questions game as in the maze-exploration game. Studies 2b and 3 ruled out alternative explanations, demonstrating that merely providing children with a visual conceptual aid to support their representation of the hypothesis space (Study 2b, $N$ = 46, 4- to 6-year-olds), or familiarising them with the hypothesis-space structure (Study 3, $N$ = 54, 4- to 6-year-olds) was not sufficient to improve their search strategies. Together, these results suggest that young children's difficulties in the 20-questions game are mainly driven by the verbal requirements of the task. However, they also demonstrate that efficient search strategies emerge much earlier than previously assumed in tasks that do not rely on verbal question generation. These findings highlight the importance of developing age-appropriate paradigms that capture children's early competence, in order to gain a more comprehensive picture of their emerging information-search abilities.

## 2.5. Summary of supplemental unpublished projects

### 2.5.1. Project 3: Where is the treasure? Investigating information-search competence in 2.5- to 4-year-olds

This project (Serko et al., in preparation) investigates the efficiency of children's (2.5 to 4-year-old) predecisional information search using a simplified version of Lindow's (2021) "Finding Presents Game". Participants had to find a treasure hidden in one of three closed boxes. The boxes all had particular shapes, colours, and images on their top, and were identical but for one feature (e.g., all boxes were blue and had a flower on top, but one was square, one was round, and one was heart-shaped). To identify the target box, children received three cards, showing all available variants of each feature on their backside (e.g., the colour card showed three blue splashes of colour, the icon card showed three flowers, the shape card showed a square, a circle, and a heart), and the feature variant of the target box on their front (e.g., a blue splash of colour, a flower, and a heart). The cards were placed frontside down in front of the children. Children were allowed to flip *exactly one card* before deciding which box to open, but were given a second attempt if they flipped the wrong card on their first go (i.e., all 3 cards were shuffled and placed down again). Crucially, as the boxes differed in only one feature (e.g., shape), only one card contained information relevant for the decision. Turning over that card was therefore both *necessary and sufficient* for making an informed decision. The results show that children's performance exceeded chance level (33%), with 61% of children flipped the relevant card on their first attempt, and an additional 19.5% did so when given a second chance ($N$ = 82), suggesting that children as young as 2.5 years are capable of identifying and consulting information necessary to make an informed decision. This suggests that the lack of systematicity found in previous studies on children's predecisional information search is not primarily related to an inability to identify relevant information. Instead it may be due to other factors, such as a tendency to additionally query irrelevant information, potentially arising from a lack of an effective stopping rule (i.e., not knowing when one has gathered enough information), or a general desire to study all available options even when that is not rationally necessary or sensible.

### 2.5.2. Project 4: Sources of inefficiency in preschoolers' information search

This project (Chai et al., in preparation) explores the informativeness of 4- to 7-year old children's information-search using a non-verbal object-selection version of the 20-question game. Children were presented with 16 cards depicting animals and plants from the "Planet Apres" scenario developed by Ruggeri et al. (2016). The 16 objects were designed such that they could be grouped into separate categories at three hierarchical levels. Children ($N$ = 67) had to find out which *kind of objects* could be found on Planet Apres by inquiring about (i.e., ask about or point at) the individual objects. Yes-No feedback was provided by moving the targeted card to a "yes" or "no" pile, indicated by two cards depicting a planet or a crossed-out planet, respectively. Children were allowed to search as long as they wanted (i.e., until they indicated that they knew which objects were found on Planet Apres), but were prompted to continue their search if they provided a wrong answer. Children's search efficiency, assessed as the *number of queries* needed to reach the solution and the *average information gain* (*IG*) of their queries, was compared to the efficiency of a simulated optimal searcher (i.e., benchmark for the best possible performance) and simulated random searcher (i.e., benchmark for chance-level performance). The results show that preschoolers' performance was significantly lower than that of a random searcher (i.e., they needed more queries, and made queries with lower *IG*). However, when excluding queries made after only one hypothesis (i.e., the correct solution) was left (i.e., questions with zero *IG*), children required significantly fewer queries than the simulated random searcher. What's more, when additionally excluding all redundant queries (i.e., queries with zero *IG*) made before only one hypothesis remained, children needed as many queries as the simulated optimal searcher. These results replicate earlier findings of a general lack of search efficiency during the preschool years. Crucially however, they also suggest that a strong tendency to make unnecessary queries (i.e., queries that provide no information gain, either before or after the solution has been identified) contributes to children's search inefficiency.

### 2.5.3. Project 5: Smart, or just lucky? Inferring question-asking competence from strategy efficiency versus effectiveness

This project (Török et al., under revision) investigates how children (5- to 10 years) and adults infer the competence of others when observing how they solved a task. Previous research has shown that children evaluate the competence of others based on how effectively someone accomplished a goal (i.e., based on the *observed outcome* of an action or strategy), such as whether an agent successfully accomplishes goal, how long they take to do so, how many attempts they need, or how complex/difficult the achieved goal was. This project extends this work by looking at whether children and adults also consider the *expected efficiency* of an action or strategy, quantified in terms of expected information gain, when inferring competence levels. Across two experiments, we varied how efficiently and how effectively different agents solved a 20-questions game (Experiments 1 and 2), where the goal is to identify a target object from a set of given options by asking only yes-no questions, and a maze-exploration game (Experiment 2), where they had to find the path through a two-dimensional maze by revealing which of its passages were open or closed.

Both tasks could be solved using two alternative strategies, a more efficient constraint-seeking strategy, targeting half of the candidate hypotheses at each step of the search, or a less efficient hypothesis-scanning strategy, targeting individual candidates sequentially. Crucially, hypothesis-scanning queries are much more affected by luck as they are, essentially, identical to guessing: if the answer to a hypothesis-scanning query is yes, the solution is found immediately. However, if the answer is no - which is generally the more likely outcome when all candidate solutions are equally likely - only one candidate solution is eliminated and the search continues until the searcher finally targets the correct solution. A hypothesis-scanning approach therefore has a highly variable actual effectiveness (i.e., may result in a very short or very long search process; in our paradigm between 1 and 8 queries). In contrast, constraint-seeking queries are more immune to effects of luck and ensure a more stable effectiveness because they rule out several candidate solutions irrespective of the feedback received. Therefore, following a constraint-seeking strategy typically results in a shorter search process (in our paradigm precisely 3 queries). As these considerations illustrate, the *expected efficiency* of an action or strategy, defined in terms of expected information gain (EIG), is a more reliable indicator of competence than *observed effectiveness*, which may depend on unrelated external factors, such as luck (e.g., the likelihood of obtaining a yes or no answer) and therefore only provides confounded information about someone's competence. We therefore equate higher competence with using the more efficient constraint-seeking strategy, and lower competence to using the less efficient, but more luck-prone, hypothesis-scanning strategy.

In Experiment 1 ($N$ = 121), only adults identified a more efficient agent as more competent, and all participants erroneously attributed higher competence to agents needing fewer questions when they employed the same inefficient strategy (reflective of the same level of competence), suggesting that they failed to consider effects of luck and and relied on the observed effectiveness to attribute competence. In Experiment 2 ($N$ = 220), adults and children from about 8 years onward successfully identified the agents using the more efficient strategy as more competent when agents were contrasted directly in a binary choice task. Overall, our results suggest that observed effectiveness is a powerful cue for competence even when such an inference may not be warranted, and that the ability to make explicit competence judgements based on efficiency (expected information gain) of a strategy alone emerges around 8 years of age, although, as shown in previous work, a more implicit understanding of competence may already be present during the preschool years.

### 2.5.4. Project 6: Ecological learning in toddlers and preschoolers

This project (Ziemann et al., in preparation) assesses the ecological learning abilities of 2- and 3-year old children. Participants play two different games, in which they either have to retrieve a ball from one of four small boxes (search game) or to catch a ball coming out of one of four exits of a marble run (catch game). In the search game, the four small boxes all have a transparent top, through which children can look but not reach inside each box, and an opening with a door on the front, through which children can reach but not look inside each box. Behind the small boxes, there is a transparent marble run with one entry on the top and four exits at the bottom, each of which is connected to one small box. During an initial training phase, children observe how several white ping pong balls move through the marble run and end up in the four small boxes. In the uniform condition, the marble run is built such that all four exits are open and the balls can end up in any of the four small boxes. In the skewed condition, three out of four paths of the marble run are blocked so that the balls all end up in the same small box. During the test phase, children are asked to cover their eyes, and the experimenter places a red ping pong ball into the marble run. Children then have to retrieve the red ball. If they are ecological learners, they should peak through the transparent tops first to see where the ball is in the uniform condition, but should reach inside the target box right away in the skewed condition.

In the catch game, children first observe how several white balls move through the marble run and exit it either always through the same exit (skewed condition), or through different exits (uniform condition). Next, children are asked to catch the red ping pong ball using either a large box covering three out of four of the marble run's exits, or a small box covering only one exit. In the uniform condition, children cannot be sure where the ball will exit and should therefore select the large box to increase their chances of catching the ball. In the skewed condition, they should select the small box as they can be certain where the ball will exit the marble run.

Children play both games, each with a different condition (e.g., search game - skewed condition; catch game - uniform condition). Data collection is still ongoing, but preliminary data indicate that children ($N$ = 33, 2- to 3-year-olds) adapt their choice of action/box as predicted: In the search game, 69% of children open a small box right away in the skewed condition (compared to 53% in the uniform condition), and 47% look inside the boxes first in the uniform condition (compared to 31% in the skewed condition). In the catch game, 60% of children pick the small box in the skewed condition (compared to 29% in the uniform condition), and 70% of children pick the large box in the uniform condition (compared to 40% in the skewed condition). Hence, this project is the first to demonstrate ecological learning in children as young as 2 years old.

# 3. Discussion

This dissertation investigated whether preschool-aged children are (a) adaptive and (b) efficient active learners, who identify and implement those active learning strategies that are best suited to provide them with the sought information given the current learning situation. To do so, it introduced several novel paradigms based on the formal and computational structure of the 20-questions game, but removing the question-asking component from it, thus making these tasks much more age-appropriate. The results from the reported studies provide compelling evidence that, under the right conditions, even toddlers and preschoolers can search for information efficiently and adaptively.

In particular, the first project (Swaboda et al., 2018) demonstrated that children as young as 3 are ecological learners, able to adapt their exploratory actions to the characteristics of the learning environment: they were more likely to shake one of two boxes first in the uniform condition, where this was necessary to resolve the uncertainty as to which large box should be opened to find an egg shaker, but preferred to open a box right away in the skewed condition, where they knew where the egg was hidden and could therefore bet on a quick win. Considering that children correctly inferred how likely the egg was to be found across the four small boxes based on the frequency pattern they observed during training, these findings are in line with previous studies demonstrating that preschoolers and even infants are excellent at tracking statistical regularities (e.g., Denison, Reed et al., 2013; Kushnir et al., 2010; Waismeyer et al., 2015; Wellman et al., 2016). Most importantly, this project showed that preschoolers exploited this statistical sensitivity to decide *which exploratory action to perform*, choosing the action that promised the largest information gain in the given situation. In demonstrating that children as young as 3 are adaptive, ecological learners, this project not only adds to the growing body of literature demonstrating that young children are excellent active learners. It additionally provides evidence for a potential key mechanism underlying children's remarkable learning capacities, namely their very ability to flexibly tailor their exploratory actions to characteristics of the current task in order to maximise learning effectiveness. Preliminary results from Project 6 (Ziemann et al., in preparation) replicate and extend these findings, providing evidence for ecological learning in children as young as 2 years old.

The second project (Swaboda et al., 2022) demonstrated that children as young as 4 searched as efficiently as 7-year-olds in a spatial search task where they could make queries non-verbally (e.g., by removing cards to reveal the path through a maze or by pointing at the features they wished to query). It also showed that merely providing children with a visual representation of the hypothesis space and/or familiarising them with the hypothesis-space structure prior to the search phase was not sufficient to improve performance. Together, the results of this project suggest that even young children possess more efficient and sophisticated information-search strategies than previously documented using the 20-questions game (Herwig, 1982; Mosher & Hornsby, 1966; Ruggeri & Feufel, 2015; Ruggeri et al., 2015, 2016), and support the view that children are skilled active learners from a very young age (e.g., Gopnik & Meltzoff, 1997; Gopnik & Wellman, 2012). Moreover, they demonstrate that children's poor performance on the 20-questions game is, at least partially, driven by the fact that their emerging verbal abilities do not allow them to fully reveal their information-search abilities, and that the wealth of developmental studies employing the traditional 20-questions game likely paints an incomplete picture of the efficiency of children's information-search strategies - especially in young children, whose verbal skills are still developing. This project therefore makes important methodological and theoretical contributions, highlighting the need for a broader array of paradigms that take into account the cognitive abilities and constraints of the targeted age groups on the one hand, and providing a potential explanation for conflicting findings of early and late emerging information-search competence on the other. Further evidence for children's early information-search competence also comes from Project 3 (Serko et al., in preparation), which showed that children as young as 2.5 years are capable of identifying and consulting information necessary for making an informed decision, at least in a constrained search context.

Furthermore, this dissertation outlines the limits of children's search efficiency. Firstly, it demonstrates that general task-demands, such as reliance on linguistic skills or task complexity, directly affect the informativeness of children's exploratory actions, potentially because these factors constrain children's exploratory repertoire (e.g., the range of available questions) or because they incur additional cognitive processing costs. For instance, in Project 4 (Chai et al., in preparation), where children had to find out which kind of objects were present on planet Apres by targeting individual objects (i.e., different types of plants and animals), children may have searched sub-optimally because they may have struggled to make the necessary abstract inferences going from the remaining individual objects to the group of objects that could potentially be found on planet Apres. Similarly, in Project 2 (Swaboda et al., 2022), children required extensive assistance with the appropriate if-then inferences upon receiving feedback in the maze-exploration game. As a pilot revealed, without such assistance children often failed to make even the easiest inferences (e.g., if the right passage on the higher level was open, the left passage on the higher level must be closed), and instead continued to test individual passages even if this provided no new information gain. It is possible that in both of these tasks, the cost of identifying the most informative queries was higher than the cost associated with simply continuing to test individual objects or passages more or less randomly - which might explain children's tendency to make unnecessary queries that provided no additional information gain.

Second, although the present results suggest that children do possess an understanding of what constitutes an efficient search strategy in a given situation, they also demonstrate that children frequently fail to apply this insight when searching freely in more unconstrained settings, making more queries than rationally necessary or even searching exhaustively (i.e., studying all available options, also see Ruggeri et al., 2016). As described above, this may, at least in parts, be due to the complexity of the task at hand. On the other hand, piloting data of Project 3 (Serko et al., in preparation) showed that children also tended to study more than the necessary number of cues when they could freely decide how many cards they wished to turn over. Considering that Project 3 employed a deliberately simple information-search paradigm, it is unlikely that task complexity can account for these findings. Instead, there may be additional factors contributing to children's tendency to keep searching even when, rationally, this is unnecessary. For instance, children may find it rewarding to see confirmed what they already know (Nickerson, 1998), or may find the mere *act of searching* rewarding, an idea that is supported by preliminary data from Project 6 (Ziemann et al., in preparation), where children greatly enjoyed opening and closing the doors of the boxes simply for the sake of seeing if there was anything inside.

Alternatively, gathering additional confirmatory evidence may actually serve a purpose for young learners, who are not only attempting to solve the current task at hand, but are also still learning how to search efficiently. For example, they may still be less certain than more experienced learners about the hypothesis space, the feedback they have received, or the constancy of what is being learned, and may therefore prefer to obtain extra feedback before giving their final answer. In line with this suggestion, Legare et al. (2013) found that the preschoolers (4- to 6-year-olds) who asked more confirmatory questions performed better overall in a 20- questions game, suggesting that the confirmatory strategy, although inefficient, may indeed provide learning advantages (also see Lapidow and Walker, 2022, for a related account of how, rationally inefficient interventions, such as gathering confirming rather than disconfirming evidence, may still be an effective learning strategy in the domain of causal learning). Together, these considerations imply that, even though children may be sensitive to the information gain of different exploratory actions, information gain alone may not be a sufficient model of how children decide how to explore. Instead, other costs (e.g., in the form of cognitive processing costs), and rewards (e.g., the fun of searching or learning benefits of confirmatory queries) of different actions also seem to play a role.

What's more, maximising search efficiency in favour of achieving functional goals may not be children's only or primary aim during play and exploration, and the mere concept of searching *efficiently* (e.g., as opposed to searching for fun) may still be emerging during the preschool years. Indeed, recent accounts propose that children's play may actually be best characterised by a willingness to incur additional costs (e.g., Chu & Schulz, 2020b). To illustrate, Chu and Schulz (2020a) instructed children to get stickers from a box standing in the middle of a spiral design printed on the floor, and varied whether they asked children to "go and try get one" (instrumental condition), or to "play and try get one" (play condition). They found that children ignored the spiral design and walked in a straight line to the box in the instrumental condition, but followed the spiral, sometimes even walking along it twice, or doings so backwards, in the play condition. Similarly, they preferred a small search space (e.g., 1 drawer) when simply asked to find a ball for another game, but a larger search space (e.g., 12 drawers) when looking for the ball was introduced as part of a hide-and-seek game. Crucially though, their behaviour in the play condition, although more costly, was not necessarily random (e.g., they carefully followed the spiral, and searched the drawers sequentially rather than arbitrarily). This suggests that children's play is not irrational, but may instead be efficient with respect to the utilities they set themselves. Importantly, such willingness to set arbitrary utilities and incur additional costs may allow children to explore new problems, to gain insights and to discover unexpected information, which they may have missed if expected information gain was the sole motivator behind their exploratory choices.

## 3.1. Implications

Returning to the opening paragraph of this dissertation, what are the implications of the presented outcomes for educational practices, and the debate about the suitability of free play for educating young children? First, the finding that preschoolers are efficient and adaptive active learners generally speaks towards the idea that children can learn through play, and may do so efficiently in play-based educational settings. Second, this dissertation highlights the conditions under which children are most likely to explore efficiently: when the task at hand is adapted to their cognitive capacities and when the search space is somewhat constrained rather than entirely free. Third, although this has not been tested directly in this dissertation, findings from earlier developmental studies (e.g., van Schijndel et al., 2015) strongly suggest that children's interventions, even if designed to be informative, can only lead to learning if they target appropriate underlying hypotheses. More generally, this point highlights an important issue for learning from free play - how can one learn something one does not even know needs to be learned? Together, these considerations imply that entirely free play may not be ideal for the construction of knowledge, a conclusion that is complemented by educational research indicating that free play alone may not be sufficient to achieve objective learning outcomes and to promote school readiness (Chi, 2009; Chien et al., 2010). On the other hand, play also has a range of benefits for children's social and self-regulatory development (e.g., Nolan & Paatsch, 2018; Pyle & Alaca, 2018; Pyle & DeLuca, 2017), and preschoolers are generally highly motivated to learn, but may lose interest in more formal instructional settings, or if performing an activity simply because an adult wants them to (Hauser, 2005; Lepper & Henderlong, 2000). Play-based approaches may therefore be more suited to spark and maintain children's interest in the taught topics, which may have important learning and motivational benefits in the long run.

A potential solution to preserve the advantages of play while optimising learning outcomes, which is directly supported by the results of this dissertation, may be to provide children with structured learning environments that they are allowed to explore on their own. An advantage of such an approach would be that learning problems can be outlined more clearly than in entirely free play settings (e.g., by preparing restricted sets of materials), which may help to focus children's exploration on a particular learning goal, all while maintaining their autonomy and sense of control over the learning process. Indeed, free exploration within structured environments is at the heart of the Montessori pedagogy, which has been associated with better academic achievements, social understanding, executive function, and an increased belief in their own abilities in preschool children (Lillard et al., 2017).

Relatedly, proponents of guided play suggest that children learn best when free exploration is accompanied by an adult, who prepares the learning environment (as in a Montessori classroom) and scaffolds children's on-going exploration (e.g., Ilgaz et al., 2018; Weisberg et al., 2016; Weisberg et al., 2015). For instance, a teacher may make comments or suggestions to encourage deeper thinking (e.g., "What do you think would happen if...?"), or provide hints about other ways to explore the problem at hand (e.g., "I wonder what would happen if you try using this one?"), thereby nudging the child towards the learning goal. Importantly, as in the Montessori pedagogy, control over the learning process remains entirely with the child, in the sense that children are free to follow or reject the adult's inputs. The effectiveness of guided play, relative to direct instruction and free play, has been demonstrated both in educational (e.g., Miller, 2018; Vogt et al., 2018) and developmental studies (for a meta-analysis, see Alfieri et al., 2011), although it is difficult to disentangle how its components, structured environments and adult scaffolding, contribute to its benefits.

Nonetheless, the results from this dissertation, and findings from cognitive psychology more generally, demonstrating that children explore best in somewhat constrained settings (e.g., Cook et al., 2011; Kushnir & Gopnik, 2005; Sobel & Sommerville, 2010; van Schijndel et al., 2015) and that their interventions actually result in learning (Bonawitz et al., 2011; Lapidow & Walker, 2020; McCormack et al., 2016; L. Schulz, Gopnik et al., 2007; Shneidman et al., 2016; Sim et al., 2017; Sobel & Sommerville, 2010), strongly point towards free exploration within structured environments as a suitable approach towards early childhood education, although further research is needed to explore this proposition within more educational settings.

## 3.2. Future directions

The present dissertation leaves open several lines for future research directions. First, the non-verbal tasks introduced in this dissertation provide great potential for the development of paradigms assessing the efficiency and adaptiveness of the exploratory strategies of even younger children, and potentially even infants. For instance, future studies could adapt the presented tasks and reduce their verbal and general cognitive processing demands even further by using eye tracking or EEG measures, which may provide a much more nuanced understanding of young children's and infants' information-search abilities. Such an approach would make important contributions to our general understanding of the developmental trajectory of children's active learning strategies, and the mechanisms underlying their cognitive development more generally.

Second, although findings from cognitive psychology provide useful insights into basic learning mechanisms, the paradigms used in developmental psychology research often have limited ecological validity as they are more structured and overly simplified than comparable problems children may encounter in real life. For instance, the quantitative framework used to assess informativeness in this dissertation can only be applied in highly constrained settings where only binary queries are possible and the hypothesis space is both constrained and predefined. This setup is not particularly representative of naturalistic learning situations, where there may be a much larger range of exploratory actions children can chose from, and where the hypotheses under consideration, as well as their likelihoods, are not predefined but determined by the learners prior experience and beliefs (Jones, Swaboda et al., 2020). To gain more ecological validity and increase the transferability of findings to more applied settings, it would be useful to develop paradigms that are more representative of real-life learning situations children may encounter in and out of the preschool class room.

Third, another important research direction would be to further explore which factors, other than task-demands, contribute to developmental changes, as well as to individual differences, in children's information-search skills. These may include cognitive abilities, such as executive function or metacognition, as well as differences in personality, motivation, parenting style, socio-economic status or cultural and environmental factors. Discerning the contributions of these factors and developing targeted early intervention programmes may be a viable way to promote children's active learning skills and to foster their natural curiosity, especially in children at risk.

Finally, there are several additional lines of inquiry on children's active learning and exploration this dissertation has only marginally touched upon, such as how children come up with a hypothesis to explore in the first place, how they decide to seek help or to continue exploring, and how they decide whom to learn from and whom to trust. For instance, a growing body of work demonstrates that children's selective trust in potential informants is driven by a complicated mixture of inferences drawn from the quality of the information provided (e.g., accuracy, completeness; see Jaswal et al., 2010; Koenig & Jaswal, 2011; Pasquini et al., 2007) and the characteristics of the agent providing the information (e.g., expertise, age, familiarity, culture; see Kinzler & Spelke, 2011; Lutz & Keil, 2002; VanderBorght & Jaswal, 2009), and that there are large developmental improvements in children's understanding of the necessary characteristics for being a reliable informant over the preschool years (e.g., Einav & Robinson, 2010; Koenig & Harris, 2005; Robinson et al., 1999; Sabbagh & Baldwin, 2001). Project 5 of this dissertation (Török et al., under revision) adds to this line of research by investigating which indicators children rely on to infer someone's *problem-solving* competence (as opposed to their factual knowledge or expertise on a particular topic). Its results show that, in a 20-questions game, children up to the age of 8 tend to infer competence from the number of questions needed to find the solution, rather than from the efficiency of the strategy used - despite the latter being the more reliable indicator of competence. Although it is questionable to what extent these particular findings are generalisable beyond the 20-questions game, this study makes an important contribution to the work on children's competence evaluations and selective trust as it is one of few studies looking more specifically at differences in *how* agents solve a problem, and thereby opens up an interesting research direction, namely the question of whether children may selectively rely on more *skilled learners* as models to learn how to learn from (see also Bridgers et al., 2018).

## 3.3. Conclusion

The studies presented in this dissertation have provided compelling evidence that preschool-aged children are efficient and adaptive active learners, who identify and implement those active learning strategies that are best suited to provide them with the sought information given the current learning situation. It has also demonstrated that children are most likely to search for information efficiently when the task at hand is adapted to their cognitive capacities and when the search space is somewhat constrained. As such, it has made an important contribution to bridging the gap between conflicting findings demonstrating early active learning competence on the one hand, and late competence on the other. In addition, it has also proposed a range of novel non-verbal paradigms, which future investigations on children's active learning may build on. Finally, by providing insights into the cognitive mechanisms underlying children's learning and the conditions under which children explore most efficiently, the results from this dissertation may also inform educational practice and policy, as they suggest that free exploration within structured environments may be a useful educational approach combining the benefits of play with a focus on enhancing learning outcomes.

# Bibliography

Alfieri, L., Brooks, P. J., Aldrich, N. J. & Tenenbaum, H. R. (2011). Does discovery-based instruction enhance learning? *Journal of Educational Psychology*, *103*, 1–18. https://doi.org/10.1037/a0021017

Aslin, R. N. (2007). What's in a look? *Developmental Science*, *10*, 48–53. https://doi.org/10.1111/j.1467-7687.2007.00563.x

Bai, Y., Ladd, H. F., Muschkin, C. G. & Dodge, K. A. (2020). Long-term effects of early childhood programs through eighth grade: Do the effects fade out or grow? *Children and Youth Services Review*, *112*, 104890. https://doi.org/10.1016/j.childyouth.2020.104890

Baillargeon, R. (1987). Object permanence in 3.5-and 4.5-month-old infants. *An Introductory Reader in Developmental Psychology*, *23*(5), 103.

Baillargeon, R., Spelke, E. S. & Wasserman, S. (1985). Object permanence in five-month-old infants. *Cognition*, *20*(3), 191–208. https://doi.org/10.1016/0010-0277(85)90008-3

Berlyne, D. E. (1969). Arousal, reward and learning. *Annals of the New York Academy of Sciences*, *159*(3), 1059–1070.

Berretty, P. M., Todd, P. M. & Blythe, P. W. (1997). Categorization by elimination: A fast and frugal approach to categorization. *Proceedings of the 19th Annual Conference of the Cognitive Science Society*, 43–48.

Berretty, P. M., Todd, P. M. & Martignon, L. (1999). Categorization by elimination: Using few cues to choose.

Betsch, T., Lehmann, A., Jekel, M., Lindow, S. & Glöckner, A. (2018). Children's application of decision strategies in a compensatory environment. *Judgment & Decision Making*, *13*(6).

Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E. & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, *120*(3), 322–330. https://doi.org/10.1016/j.cognition.2010.10.001

Bonawitz, E., van Schijndel, T. J., 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

Bridgers, S., Gweon, H., Bretzke, M. & Ruggeri, A. (2018). How you learned matters: The process by which others learn informs young children's decisions about whom to ask for help. *Cognitive Science*, 1402–1407.

Brown, C. P. (2021). *Resisting the kinder-race: Restoring joy to early learning*. Teachers College Press.

Bruner, J. S., Jolly, A. & Sylva, K. (1976). *Play: Its role in development and evolution*. Penguin.

Butterworth, G. & Hopkins, B. (1988). Hand-mouth coordination in the new-born baby. *British Journal of Developmental Psychology*, *6*(4), 303–314. https://doi.org/10.1111/j.2044-835X.1988.tb01103.x

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

Carey, S. (1985). *Conceptual change in childhood*. MIT Press.

Chai, K., Xu, F., Swaboda, N. & Ruggeri, A. (in preparation). Sources of inefficiency in preschoolers' information search.

Chi, M. T. (2009). Active-constructive-interactive: A conceptual framework for differentiating learning activities. *Topics in Cognitive Science*, *1*(1), 73–105. https://doi.org/10.1111/j.1756-8765.2008.01005.x

Chien, N. C., Howes, C., Burchinal, M., Pianta, R. C., Ritchie, S., Bryant, D. M., Clifford, R. M., Early, D. M. & Barbarin, O. A. (2010). Children's classroom engagement and school readiness gains in prekindergarten. *Child Development*, *81*(5), 1534–1549. https://doi.org/10.1111/j.1467-8624.2010.01490.x

Chin, J., Payne, B. R., Fu, W.-T., Morrow, D. G. & Stine-Morrow, E. A. (2015). Information foraging across the life span: Search and switch in unknown patches. *Topics in Cognitive Science*, *7*(3), 428–450. https://doi.org/10.1111/tops.12147

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.

Chu, J. & Schulz, L. (2020a). Exploratory play, rational action, and efficient search. https://doi.org/10.31234/osf.io/9yra2

Chu, J. & Schulz, L. (2020b). Play, curiosity, and cognition. *Annual Review of Developmental Psychology*, *2*, 317–343. https://doi.org/10.1146/annurev-devpsych-070120-014806

Coenen, A., Ruggeri, A., Bramley, N. R. & Gureckis, T. M. (2019). Testing one or multiple: How beliefs about sparsity affect causal experimentation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *45*(11), 1923. https://doi.org/10.1037/xlm0000680

Cook, C., Goodman, N. D. & Schulz, L. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, *120*(3), 341–349. https://doi.org/10.1016/j.cognition.2011.03.003

Davidson, D. (1991a). Children's decision-making examined with an information-board procedure. *Cognitive Development*, *6*(1), 77–90. https://doi.org/10.1016/0885-2014(91)90007-Z

Davidson, D. (1991b). Developmental differences in children's search of predecisional information. *Journal of Experimental Child Psychology*, *52*(2), 239–255. https://doi.org/10.1016/0022-0965(91)90061-V

Davidson, D. (1996). The effects of decision characteristics on children's selective search of predecisional information. *Acta Psychologica*, *92*(3), 263–281. https://doi.org/10.1016/0001-6918(95)00014-3

Denison, S., Bonawitz, E., Gopnik, A. & Griffiths, T. L. (2013). Rational variability in children's causal inferences: The sampling hypothesis. *Cognition*, *126*(2), 285–300. https://doi.org/10.1016/j.cognition.2012.10.010

Denison, S., Reed, C. & Xu, F. (2013). The emergence of probabilistic reasoning in very young infants: Evidence from 4.5-and 6-month-olds. *Developmental Psychology*, *49*(2), 243. https://doi.org/10.1037/a0028278

Denison, S. & Xu, F. (2014). The origins of probabilistic inference in human infants. *Cognition*, *130*(3), 335–347. https://doi.org/10.1016/j.cognition.2013.12.001

Denison, S. & Xu, F. (2019). Infant statisticians: The origins of reasoning under uncertainty. *Perspectives on Psychological Science*, *14*(4), 499–509. https://doi.org/10.1177/1745691619847201

Denney, D. R. & Denney, N. W. (1973). The use of classification for problem solving: A comparison of middle and old age. *Developmental Psychology*, *9*(2), 275. https://doi.org/10.1037/h0035092

Dewey, J. (1986). Experience and education. *The Educational Forum*, *50*(3), 241–252. https://doi.org/10.1080/00131728609335764

Duh, S. & Wang, S.-h. (2019). Infants detect patterns of choices despite counter evidence, but timing of inconsistency matters. *Journal of Cognition and Development*, *20*(1), 96–106. https://doi.org/10.1080/15248372.2018.1528976

Eaves Jr, B. S. & Shafto, P. (2012). Unifying pedagogical reasoning and epistemic trust. *Advances in Child Development and Behavior*, *43*, 295–319. https://doi.org/10.1016/B978-0-12-397919-3.00011-3

Eimas, P. D. (1970). Information processing in problem solving as a function of developmental level and stimulus saliency. *Developmental Psychology*, *2*(2), 224. https://doi.org/10.1037/h0028746

Einav, S. & Robinson, E. J. (2010). Children's sensitivity to error magnitude when evaluating informants. *Cognitive Development*, *25*(3), 218–232. https://doi.org/10.1016/j.cogdev.2010.04.002

Faas, S., Wu, S.-c. & Geiger, S. (2017). The importance of play in early childhood education: A critical perspective on current policies and practices in germany and hong kong. *Global Education Review*, *4*(2).

Fagan, J. F. (1970). Memory in the infant. *Journal of Experimental Child Psychology*, *9*(2), 217–226. https://doi.org/10.1016/0022-0965(70)90087-1

Fantz, R. L. (1964). Visual experience in infants: Decreased attention to familiar patterns relative to novel ones. *Science*, *146*(3644), 668–670. https://doi.org/10.1126/science.146.3644.668

Fiser, J. & Aslin, R. N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences*, *99*(24), 15822–15826. https://doi.org/10.1073/pnas.232472899

Fontanari, L., Gonzalez, M., Vallortigara, G. & Girotto, V. (2014). Probabilistic cognition in two indigenous mayan groups. *Proceedings of the National Academy of Sciences*, *111*(48), 17075–17080. https://doi.org/10.1073/pnas.1410583111

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. https://doi.org/10.1111/j.1467-8624.2009.01356.x

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. https://doi.org/10.1080/15248372.2015.1098649

Fröbel, F. (1839). Die bildung der kinder vor dem schulfähigen alter, und die ausführung einer bildungsanstalt zu erziehern und pflegern in dem angegebenen alter, besonders die bildung zu lehrern an kleinkinderschulen betreffend. *Kommt lasst uns unseren Kindern leben*, 184–188.

Fröbel, F. (1885). *The education of man*. New York: A. Lovell Company.

Girotto, V. & Gonzalez, M. (2008). Children's understanding of posterior probability. *Cognition*, *106*(1), 325–344. https://doi.org/10.1016/j.cognition.2007.02.005

Gonzalez, M. & Girotto, V. (2011). Combinatorics and probability: Six-to ten-year-olds reliably predict whether a relation will occur. *Cognition*, *120*(3), 372–379. https://doi.org/10.1016/j.cognition.2010.10.006

Gopnik, A. & Meltzoff, A. N. (1997). *Words, thoughts, and theories*. MIT Press.

Gopnik, A., Sobel, D. M., Schulz, L. & 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. https://doi.org/10.1037/0012-1649.37.5.620

Gopnik, A. & Wellman, H. M. (1992). Why the child's theory of mind really is a theory.

Gopnik, A. & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, bayesian learning mechanisms, and the theory theory. *Psychological Bulletin*, *138*(6), 1085. https://doi.org/10.1037/a0028044

Gottlieb, J., Oudeyer, P.-Y., Lopes, M. & Baranes, A. (2013). Information-seeking, curiosity, and attention: Computational and neural mechanisms. *Trends in Cognitive Sciences*, *17*(11), 585–593. https://doi.org/10.1016/j.tics.2013.09.001

Green, L. & Mehr, D. R. (1997). What alters physicians' decisions to admit to the coronary care unit? *Journal of Family Practice*, *45*(3), 219–226.

Gregan-Paxton, J. & John, D. R. (1995). Are young children adaptive decision makers? a study of age differences in information search behavior. *Journal of Consumer Research*, *21*(4), 567–580. https://doi.org/10.1086/209419

Gregan-Paxton, J. & John, D. R. (1997). The emergence of adaptive decision making in children. *Journal of Consumer Research*, *24*(1), 43–56. https://doi.org/10.1086/209492

Greif, M. L., Kemler Nelson, D. G., Keil, F. C. & Gutierrez, F. (2006). What do children want to know about animals and artifacts? domain-specific requests for information. *Psychological Science*, *17*(6), 455–459. https://doi.org/10.1111/j.1467-9280.2006.01727.x

Gweon, H. & Schulz, L. (2008). Stretching to learn: Ambiguous evidence and variability in preschoolers' exploratory play. *Proceedings of the 30th Annual Meeting of the Cognitive Science Society*, 570–574.

Hamilton, G. C. (2003). *Emergency medicine: An approach to clinical problem-solving*. Elsevier Health Sciences. https://doi.org/10.1016/j.jemermed.2003.10.007

Hamlin, K. J., Wynn, K. & Bloom, P. (2010). Three-month-olds show a negativity bias in their social evaluations. *Developmental Science*, *13*(6), 923–929. https://doi.org/10.1111/j.1467-7687.2010.00951.x

Hanus, D. & Call, J. (2014). When maths trumps logic: Probabilistic judgements in chimpanzees. *Biology Letters*, *10*(12), 20140892. https://doi.org/10.1098/rsbl.2014.0892

Harris, P. L. (2000). On not falling down to earth: Children's metaphysical questions. *Imagining the Impossible*, 157–178.

Harris, P. L. (2012). *Trusting what you're told: How children learn from others*. Cambridge, MA: The Belknap Press/Harvard University Press. https://doi.org/10.4159/harvard.9780674065192

Harris, P. L. & Koenig, M. A. (2006). Trust in testimony: How children learn about science and religion. *Child Development*, *77*(3), 505–524. https://doi.org/10.1111/j.1467-8624.2006.00886.x

Hauser, B. (2005). Das spiel als lernmodus: Unter druck von verschulung–im lichte der neueren forschung. *Bildung*, *4*, 143–167.

Herwig, J. E. (1982). Effects of age, stimuli, and category recognition factors in children's inquiry behavior. *Journal of Experimental Child Psychology*, *33*(2), 196–206. https://doi.org/10.1016/0022-0965(82)90015-7

Hickling, A. K. & Wellman, H. M. (2001). The emergence of children's causal explanations and theories: Evidence from everyday conversation. *Developmental Psychology*, *37*(5), 668. https://doi.org/10.1037/0012-1649.37.5.668

Hirsh-Pasek, K., Hadani, H., Blinkoff, E. & Golinkoff, R. (2020). A new path to education reform: Playful learning promotes 21st century skills in school and beyond. *Policy Brief*.

Howse, R. B., Best, D. L. & Stone, E. R. (2003). Children's decision making: The effects of training, reinforcement, and memory aids. *Cognitive Development*, *18*(2), 247–268. https://doi.org/10.1016/S0885-2014(03)00023-6

Ilgaz, H., Hassinger-Das, B., Hirsh-Pasek, K. & Golinkoff, R. M. (2018). Making the case for playful learning. *International handbook of early childhood education* (pp. 1245–1263). Springer. https://doi.org/10.1007/978-94-024-0927-7

Jaswal, V. K., Croft, A. C., Setia, A. R. & Cole, C. A. (2010). Young children have a specific, highly robust bias to trust testimony. *Psychological Science*, *21*(10), 1541–1547. https://doi.org/10.1177/0956797610383438

Jones, A., Bramley, N., Gureckis, T. & Ruggeri, A. (2020). Changing many things at once sometimes makes for a good experiment, and children know that. https://doi.org/10.31234/osf.io/9qv5y

Jones, A., Swaboda, N. & Ruggeri, A. (2020). Developmental changes in question-asking. In L. P. Butler, S. Ronfard & K. H. Corriveau (Eds.), *The questioning child: Insights from psychology and education* (pp. 118–143). Cambridge University Press. https://doi.org/10.1017/9781108553803.007

Kachergis, G., Rhodes, M. & Gureckis, T. (2017). Desirable difficulties during the development of active inquiry skills. *Cognition*, *166*, 407–417. https://doi.org/10.1016/j.cognition.2017.05.021

Kaldy, Z. & Blaser, E. (2013). Red to green or fast to slow? infants' visual working memory for "just salient differences". *Child Development*, *84*(6), 1855–1862. https://doi.org/10.1111/cdev.12086

Kelemen, D., Callanan, M. A., Casler, K. & Pérez-Granados, D. R. (2005). Why things happen: Teleological explanation in parent-child conversations. *Developmental Psychology*, *41*(1), 251. https://doi.org/10.1037/0012-1649.41.1.251

Kemler, D. G., Nelson, Louisa Chan, E. & Holt, M. B. (2004). When children ask, "what is it?" what do they want to know about artifacts? *Psychological Science*, *15*(6), 384–389. https://doi.org/10.1111/j.0956-7976.2004.00689.x

Kidd, C. & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, *88*(3), 449–460. https://doi.org/10.1016/j.neuron.2015.09.010

Kidd, C., Piantadosi, S. T. & Aslin, R. N. (2012). The goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PloS One*, *7*(5), e36399. https://doi.org/10.1371/journal.pone.0036399

Kidd, C., Piantadosi, S. T. & Aslin, R. N. (2014). The goldilocks effect in infant auditory attention. *Child Development*, *85*(5), 1795–1804. https://doi.org/10.1111/j.1467-8624.2005.00849.x

Kim, I. K. & Spelke, E. S. (1992). Infants' sensitivity to effects of gravity on visible object motion. *Journal of Experimental Psychology: Human Perception and Performance*, *18*(2), 385.

Kim, I. K. & Spelke, E. S. (1999). Perception and understanding of effects of gravity and inertia on object motion. *Developmental Science*, *2*(3), 339–362. https://doi.org/10.1111/1467-7687.00080

Kinzler, K. D. & Spelke, E. S. (2011). Do infants show social preferences for people differing in race? *Cognition*, *119*(1), 1–9. https://doi.org/10.1016/j.cognition.2010.10.019

Kirkham, N. Z., Slemmer, J. A. & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, *83*(2), B35–B42. https://doi.org/10.1016/S0010-0277(02)00004-5

Klahr, D., Fay, A. L. & Dunbar, K. (1993). Heuristics for scientific experimentation: A developmental study. *Cognitive Psychology*, *25*, 111–146. https://doi.org/10.1006/cogp.1993.1003

Koenig, M. A. & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, *76*(6), 1261–1277. https://doi.org/10.1111/j.1467-8624.2005.00849.x

Koenig, M. A. & Jaswal, V. K. (2011). Characterizing children's expectations about expertise and incompetence: Halo or pitchfork effects? *Child Development*, *82*(5), 1634–1647. https://doi.org/10.1111/j.1467-8624.2011.01618.x

Köksal, Ö., Sodian, B. & Legare, C. H. (2021). Young children's metacognitive awareness of confounded evidence. *Journal of Experimental Child Psychology*, *205*, 105080. https://doi.org/10.1016/j.jecp.2020.105080

Kuhlmeier, V., Wynn, K. & Bloom, P. (2003). Attribution of dispositional states by 12-month-olds. *Psychological Science*, *14*(5), 402–408. https://doi.org/10.1111/1467-9280.01454

Kurkul, K. E. & Corriveau, K. H. (2018). Question, explanation, follow-up: A mechanism for learning from others? *Child Development*, *89*(1), 280–294. https://doi.org/10.1111/cdev.12726

Kushnir, T. & Gopnik, A. (2005). Young children infer causal strength from probabilities and interventions. *Psychological Science*, *16*(9), 678–683. https://doi.org/10.1111/j.1467-9280.2005.01595.x

Kushnir, T., Xu, F. & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological Science*, *21*(8), 1134–1140. https://doi.org/10.1177/0956797610376652

Lang, A. & Betsch, T. (2018). Children's neglect of probabilities in decision making with and without feedback. *Frontiers in Psychology*, *9*, 191. https://doi.org/10.3389/fpsyg.2018.00191

Lapidow, E. & Walker, C. M. (2020). Informative experimentation in intuitive science: Children select and learn from their own causal interventions. *Cognition*, *201*, 104315. https://doi.org/10.1016/j.cognition.2020.104315

Lapidow, E. & Walker, C. M. (2022). Rethinking the "gap": Self-directed learning in cognitive development and scientific reasoning. *Wiley Interdisciplinary Reviews: Cognitive Science*, *13*(2), e1580. https://doi.org/10.1002/wcs.1580

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., Mills, C. M., Souza, A. L., Plummer, L. E. & Yasskin, R. (2013). The use of questions as problem-solving strategies during early childhood. *Journal of Experimental Child Psychology*, *114*(1), 63–76. https://doi.org/10.1016/j.jecp.2012.07.002

Lepper, M. R. & Henderlong, J. (2000). Turning “play” into “work” and “work” into “play”: 25 years of research on intrinsic versus extrinsic motivation. *Intrinsic and Extrinsic Motivation*, 257–307. https://doi.org/10.1016/B978-012619070-0/50032-5

Lillard, A. S., Heise, M. J., Richey, E. M., Tong, X., Hart, A. & Bray, P. M. (2017). Montessori preschool elevates and equalizes child outcomes: A longitudinal study. *Frontiers in Psychology*, 1783.

Lindow, S. (2021). On searching and finding: The development of information search abilities. *Cognitive Development*, *58*, 101011. https://doi.org/10.1016/j.cogdev.2021.101011

Lindow, S. & Betsch, T. (2019). Children’s adaptive decision making and the costs of information search. *Journal of Applied Developmental Psychology*, *60*, 24–34. https://doi.org/10.1016/j.appdev.2018.09.006

Liu, S., Ullman, T. D., Tenenbaum, J. B. & Spelke, E. S. (2017). Ten-month-old infants infer the value of goals from the costs of actions. *Science*, *358*(6366), 1038–1041. https://doi.org/10.1126/science.aag2132

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, *116*(1), 75. https://doi.org/10.1037/0033-2909.116.1.75

Luo, Y. (2011). Three-month-old infants attribute goals to a non-human agent. *Developmental Science*, *14*(2), 453–460. https://doi.org/10.1111/j.1467-7687.2010.00995.x

Lutz, D. J. & Keil, F. C. (2002). Early understanding of the division of cognitive labor. *Child Development*, *73*(4), 1073–1084. https://doi.org/10.1111/1467-8624.00458

Magnuson, K. A., Ruhm, C. & Waldfogel, J. (2007). Does prekindergarten improve school preparation and performance? [The Economics of Early Childhood Education]. *Economics of Education Review*, *26*(1), 33–51. https://doi.org/10.1016/j.econedurev.2005.09.008

Martignon, L., Katsikopoulos, K. V. & Woike, J. K. (2008). Categorization with limited resources: A family of simple heuristics. *Journal of Mathematical Psychology*, *52*(6), 352–361. https://doi.org/10.1016/j.jmp.2008.04.003

McCormack, T., Bramley, N., Frosch, C., Patrick, F. & Lagnado, D. (2016). Children's use of interventions to learn causal structure. *Journal of Experimental Child Psychology*, *141*, 1–22. https://doi.org/10.1016/j.jecp.2015.06.017

McCoy, D. C., Yoshikawa, H., Ziol-Guest, K. M., Duncan, G. J., Schindler, H. S., Magnuson, K., Yang, R., Koepp, A. & Shonkoff, J. P. (2017). Impacts of early childhood education on medium-and long-term educational outcomes. *Educational Researcher*, *46*(8), 474–487. https://doi.org/10.3102/0013189X17737739

McCrink, K. & Wynn, K. (2004). Large-number addition and subtraction by 9-month-old infants. *Psychological Science*, *15*(11), 776–781. https://doi.org/10.1111/j.0956-7976.2004.00755.x

Meder, B., Nelson, J. D., Jones, M. & Ruggeri, A. (2019). Stepwise versus globally optimal search in children and adults. *Cognition*, *191*, 103965. https://doi.org/10.1016/j.cognition.2019.05.002

Miller, T. (2018). Developing numeracy skills using interactive technology in a play-based learning environment. *International Journal of STEM Education*, *5*(1), 39. https://doi.org/10.1186/s40594-018-0135-2

Mills, C. M. (2013). Knowing when to doubt: Developing a critical stance when learning from others. *Developmental Psychology*, *49*(3), 404. https://doi.org/10.1037/a0029500

Mills, C. M., Legare, C. H., Bills, M. & Mejias, C. (2010). Preschoolers use questions as a tool to acquire knowledge from different sources. *Journal of Cognition and Development*, *11*(4), 533–560. https://doi.org/10.1080/15248372.2010.516419

Moeller, A., Sodian, B. & Sobel, D. M. (2022). Developmental trajectories in diagnostic reasoning: Understanding data are confounded develops independently of choosing informative interventions to resolve confounded data. *Frontiers in Psychology*, *13*, 800226–800226. https://doi.org/10.3389/fpsyg.2022.800226

Montessori, M. (1912). The montessori method. *Schocken. New York*.

Montessori, M. (1967). *The absorbent mind*. New York, NY: Holt, Rinehart  Winston.

Mosher, F. & Hornsby, J. (1966). On asking questions. In J. Bruner, R. Oliver & P. Greenfield (Eds.), *Studies in cognitive growth*. Wiley New York, NY.

Muentener, P. & Schulz, L. (2014). Toddlers infer unobserved causes for spontaneous events. *Frontiers in Psychology*, *5*, 1496. https://doi.org/10.3389/fpsyg.2014.01496

Navarro, D. J. & Perfors, A. F. (2011). Hypothesis generation, sparse categories, and the positive test strategy. *Psychological Review*, *118*(1), 120–135. https://doi.org/10.1037/a0021110

Nelson, J. D., Divjak, B., Gudmundsdottir, G., Martignon, L. F. & Meder, B. (2014). Children's sequential information search is sensitive to environmental probabilities. *Cognition*, *130*(1), 74–80. https://doi.org/10.1016/j.cognition.2013.09.007

Nelson, J. D., McKenzie, C. R., Cottrell, G. W. & Sejnowski, T. J. (2010). Experience matters: Information acquisition optimizes probability gain. *Psychological Science*, *21*(7), 960–969. https://doi.org/10.1177/0956797610372637

Nelson, J. D., Meder, B. & Jones, M. (2018). Towards a theory of heuristic and optimal planning for sequential information search. *PsyArXiv*. https://doi.org/10.31234/osf.io/bxdf4

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, *2*(2), 175–220. https://doi.org/10.1037/1089-2680.2.2.175

Nicolopoulou, A. (2010). The alarming disappearance of play from early childhood education. *Human Development*, *53*(1), 1–4. https://doi.org/10.1159/000268135

Nolan, A. & Paatsch, L. (2018). (re) affirming identities: Implementing a play-based approach to learning in the early years of schooling. *International Journal of Early Years Education*, *26*(1), 42–55. https://doi.org/10.1080/09669760.2017.1369397

Notaro, P. C., Gelman, S. A. & Zimmerman, M. A. (2002). Biases in reasoning about the consequences of psychogenic bodily reactions: Domain boundaries in cognitive development. *Merrill-Palmer Quarterly (1982-)*, 427–449.

Oakes, L. M. (2010). Using habituation of looking time to assess mental processes in infancy. *Journal of Cognition and Development*, *11*(3), 255–268. https://doi.org/10.1080/15248371003699977

Oakes, L. M. (2017). Sample size, statistical power, and false conclusions in infant looking-time research. *Infancy*, *22*(4), 436–469. https://doi.org/0.1111/infa.12186

Oaksford, M. & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, *101*(4), 608.

Onishi, K. H. & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, *308*(5719), 255–258. https://doi.org/10.1126/science.1107621

Oudeyer, P.-Y. (2018). Computational theories of curiosity-driven learning. *The new science of curiosity*. https://doi.org/10.48550/arXiv.1802.10546

Pasquini, E. S., Corriveau, K. H., Koenig, M. & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental psychology*, *43*(5), 1216. https://doi.org/10.1037/0012-1649.43.5.1216

Perez, J. & Feigenson, L. (2022). Violations of expectation trigger infants to search for explanations. *Cognition*, *218*, 104942. https://doi.org/10.1016/j.cognition.2021.104942

Piaget, J. (1929). *The child's conception of the world*. London: Routledge.

Piaget, J. (1954). *The construction of reality in the child*. Trans. M. Cook, Basic Books, New York. https://doi.org/10.1037/11168-000

Powell, L. J. & Spelke, E. S. (2013). Preverbal infants expect members of social groups to act alike. *Proceedings of the National Academy of Sciences*, *110*(41), E3965–E3972. https://doi.org/10.1073/pnas.1304326110

Pyle, A. & Alaca, B. (2018). Kindergarten children's perspectives on play and learning. *Early Child Development and Care*, *188*(8), 1063–1075. https://doi.org/10.1080/03004430.2016.1245190

Pyle, A. & DeLuca, C. (2017). Assessment in play-based kindergarten classrooms: An empirical study of teacher practice. https://doi.org/10.1080/00220671.2015.1118005

Rakoczy, H., Clüver, A., Saucke, L., Stoffregen, N., Gräbener, A., Migura, J. & Call, J. (2014). Apes are intuitive statisticians. *Cognition*, *131*(1), 60–68. https://doi.org/10.1016/j.cognition.2013.12.011

Robinson, E., Champion, H. & Mitchell, P. (1999). Children's ability to infer utterance veracity from speaker informedness. *Developmental Psychology*, *35*(2), 535. https://doi.org/10.1037//0012-1649.35.2.535

Rochat, P. (1983). Oral touch in young infants: Response to variations of nipple characteristics in the first months of life. *International Journal of Behavioral Development*, *6*(2), 123–133. https://doi.org/10.1177/016502548300600201

Rochat, P. (1987). Mouthing and grasping in neonates: Evidence for the early detection of what hard or soft substances afford for action. *Infant Behavior and Development*, *10*(4), 435–449. https://doi.org/10.1016/0163-6383(87)90041-5

Rochat, P. (1989). Object manipulation and exploration in 2-to 5-month-old infants. *Developmental Psychology*, *25*(6), 871. https://doi.org/10.1037/0012-1649.25.6.871

Ronfard, S., Zambrana, I. M., Hermansen, T. K. & Kelemen, D. (2018). Question-asking in childhood: A review of the literature and a framework for understanding its development. *Developmental Review*, *49*, 101–120. https://doi.org/10.1016/j.dr.2018.05.002

Roßbach, H.-G. (2008). Was und wie sollen kinder im kindergarten lernen? *Die andere seite der bildung* (pp. 123–131). Springer.

Ruggeri, A. & Feufel, M. (2015). How basic-level objects facilitate question-asking in a categorization task. *Frontiers in Psychology*, *6*, 918. https://doi.org/10.3389/fpsyg.2015.00918

Ruggeri, A. & Lombrozo, T. (2014). Learning by asking: How children ask questions to achieve efficient search. *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, *36*(36).

Ruggeri, A., Lombrozo, T., Griffiths, T. L. & Xu, F. (2015). Children search for information as efficiently as adults, but seek additional confirmatory evidence. *Proceedings of the 37th Annual Meeting of the Cognitive Science Society*.

Ruggeri, A., Lombrozo, T., Griffiths, T. L. & Xu, F. (2016). Sources of developmental change in the efficiency of information search. *Developmental Psychology*, *52*(12), 2159. https://doi.org/10.1037/dev0000240

Ruggeri, A., Sim, Z. L. & Xu, F. (2017). “Why is Toma late to school again?" Preschoolers identify the most informative questions. *Developmental Psychology*, *53*(9), 1620. https://doi.org/10.1037/dev0000340

Ruggeri, A., Swaboda, N., Sim, Z. L. & Gopnik, A. (2019). Shake it baby, but only when needed: Preschoolers adapt their exploratory strategies to the information structure of the task. *Cognition*, *193*, 104013. https://doi.org/10.1016/j.cognition.2019.104013

Ruggeri, A., Walker, C. M., Lombrozo, T. & Gopnik, A. (2021). How to help young children ask better questions? *Frontiers in Psychology*, *11*, 2908. https://doi.org/10.3389/fpsyg.2020.586819

Saayman, G., Ames, E. W. & Moffett, A. (1964). Response to novelty as an indicator of visual discrimination in the human infant. *Journal of Experimental Child Psychology*, *1*(2), 189–198. https://doi.org/10.1016/0022-0965(64)90021-9

Sabbagh, M. A. & Baldwin, D. A. (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers’ theory of mind and semantic development. *Child Development*, *72*(4), 1054–1070. https://doi.org/10.1111/1467-8624.00334

Saffran, J. R., Aslin, R. N. & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*(5294), 1926–1928. https://doi.org/10.1126/science.274.5294.1926

Sahlberg, P. & Doyle, W. (2019). *Let the children play: How more play will save our schools and help children thrive*. Oxford University Press, USA.

Schauble, L., Klopfer, L. E. & Raghavan, K. (1991). Students’ transition from an engineering model to a science model of experimentation. *Journal of Research in Science Teaching*, *28*(9), 859–882. https://doi.org/10.1002/tea.3660280910

Schulz, E., Wu, C. M., Ruggeri, A. & Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychological Science*, *30*(11), 1561–1572. https://doi.org/10.1177/0956797619863663

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. & Griffiths, T. L. (2007). Can being scared cause tummy aches? naive theories, ambiguous evidence, and preschoolers' causal inferences. *Developmental Psychology*, *43*(5), 1124. https://doi.org/10.1037/0012-1649.43.5.1124

Schulz, L., Gopnik, A. & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science*, *10*(3), 322–332. https://doi.org/10.1111/j.1467-7687.2007.00587.x

Serko, D., Swaboda, N. & Ruggeri, A. (in preparation). Where is the treasure? investigating information-search competence in 2.5- to 4-year-olds.

Shannon, C. E. (1948). A mathematical theory of communication, part i, part ii. *Bell Syst. Tech. J.*, *27*, 623–656. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

Shneidman, L., Gweon, H., Schulz, L. & Woodward, A. L. (2016). Learning from others and spontaneous exploration: A cross-cultural investigation. *Child Development*, *87*(3), 723–735. https://doi.org/10.1111/cdev.12502

Shtulman, A. & Walker, C. (2020). Developing an understanding of science. *Annual Review of Developmental Psychology*, *2*(1), 111–132. https://doi.org/10.1146/annurev-devpsych-060320-092346

Siegel, M. H., Magid, R. W., Pelz, M., Tenenbaum, J. B. & Schulz, L. (2021). Children's exploratory play tracks the discriminability of hypotheses. *Nature Communications*, *12*(1), 1–9. https://doi.org/10.1038/s41467-021-23431-2

Siegler, R. S. (1977). The twenty questions game as a form of problem solving. *Child Development*, 395–403. https://doi.org/10.2307/1128632

Sim, Z. L., Mahal, K. & Xu, F. (2017). Learning about causal systems through play. *Proceedings of the 39th Annual Conference of the Cognitive Science Society*.

Sim, Z. L. & Xu, F. (2017a). Infants preferentially approach and explore the unexpected. *British Journal of Developmental Psychology*, *35*(4), 596–608. https://doi.org/10.1111/bjdp.12198

Sim, Z. L. & Xu, F. (2017b). Learning higher-order generalizations through free play: Evidence from 2-and 3-year-old children. *Developmental Psychology*, *53*(4), 642. https://doi.org/10.1037/dev0000278

Skerry, A. E. & Spelke, E. S. (2014). Preverbal infants identify emotional reactions that are incongruent with goal outcomes. *Cognition*, *130*(2), 204–216. https://doi.org/10.1016/j.cognition.2013.11.002

Sobel, D. M. & Kushnir, T. (2013). Knowledge matters: How children evaluate the reliability of testimony as a process of rational inference. *Psychological Review*, *120*(4), 779. https://doi.org/10.1037/a0034191

Sobel, D. M. & Sommerville, J. (2010). The importance of discovery in children's causal learning from interventions. *Frontiers in Psychology*, *1*, 176. https://doi.org/10.3389/fpsyg.2010.00176

Spelke, E. S., Breinlinger, K., Macomber, J. & Jacobson, K. (1992). Origins of knowledge. *Psychological Review*, *99*(4), 605. https://doi.org/10.1037/0033-295X.99.4.605

Spelke, E. S. & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, *10*(1), 89–96. https://doi.org/10.1111/j.1467-7687.2007.00569.x

Stahl, A. E. & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, *348*(6230), 91–94. https://doi.org/10.1126/science.aaa3799

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E.-J. & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*(3), 453–489. https://doi.org/10.1207/s15516709cog2703_6

Swaboda, N., Meder, B. & Ruggeri, A. (2022). Finding the (most efficient) way out of a maze is easier than asking (good) questions. *Developmental Psychology*, Advance online publication. https://doi.org/10.1037/dev0001386

Swaboda, N., Ruggeri, A. & Gopnik, A. (2018). Preschoolers adapt their exploratory strategies to the information structure of the task. *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*, 1102–1107. https://cogsci.mindmodeling.org/2018/papers/0218/0218.pdf

Tecwyn, E. C., Denison, S., Messer, E. J. & Buchsbaum, D. (2017). Intuitive probabilistic inference in capuchin monkeys. *Animal Cognition*, *20*(2), 243–256. https://doi.org/10.1007/s10071-016-1043-9

Thornton, S. (1982). Challenging "early competence": A process oriented analysis of children's classifying. *Cognitive Science*, *6*(1), 77–100. https://doi.org/10.1016/S0364-0213(82)80006-2

Todd, P. M. & Gigerenzer, G. E. (2012). *Ecological rationality: Intelligence in the world*. Oxford University Press.

Török, G., Swaboda, N. & Ruggeri, A. (under revision). Smart, or just lucky? inferring question-asking competence from strategies' efficiency versus effectiveness.

VanderBorght, M. & Jaswal, V. K. (2009). Who knows best? preschoolers sometimes prefer child informants over adult informants. *Infant and Child Development: An International Journal of Research and Practice*, *18*(1), 61–71. https://doi.org/10.1002/icd.591

van Schijndel, T. J., Visser, I., van Bers, B. M. & Raijmakers, M. E. (2015). Preschoolers perform more informative experiments after observing theory-violating evidence. *Journal of Experimental Child Psychology*, *131*, 104–119. https://doi.org/10.1016/j.jecp.2014.11.008

Vogt, F., Hauser, B., Stebler, R., Rechsteiner, K. & Urech, C. (2018). Learning through play–pedagogy and learning outcomes in early childhood mathematics. *European Early Childhood Education Research Journal*, *26*(4), 589–603. https://doi.org/10.1080/1350293X.2018.1487160

Vygotsky, L. S. & Wollock, J. (1997). *The collected works of ls vygotsky: Problems of the theory and history of psychology (vol. 3).* Springer Science & Business Media New York.

Waismeyer, A., Meltzoff, A. N. & Gopnik, A. (2015). Causal learning from probabilistic events in 24-month-olds: An action measure. *Developmental Science*, *18*(1), 175–182. https://doi.org/10.1111/desc.12208

Wason, P. (1960). On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, *12(3)*, 129–140. https://doi.org/10.1080/17470216008416717

Weisberg, D. S., Hirsh-Pasek, K., Golinkoff, R. M., Kittredge, A. K. & Klahr, D. (2016). Guided play: Principles and practices. *Current Directions in Psychological Science*, *25*(3), 177–182. https://doi.org/10.1177/0963721416645512

Weisberg, D. S., Kittredge, A. K., Hirsh-Pasek, K., Golinkoff, R. M. & Klahr, D. (2015). Making play work for education. *Phi Delta Kappan*, *96*(8), 8–13.

Wellman, H. M. & Gelman, S. A. (1992). Cognitive development: Foundational theories of core domains. *Annual Review of Psychology*, *43*(1), 337–375.

Wellman, H. M., Kushnir, T., Xu, F. & Brink, K. A. (2016). Infants use statistical sampling to understand the psychological world. *Infancy*, *21*(5), 668–676. https://doi.org/10.1111/infa.12131

Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, *69*(1), 1–34. https://doi.org/10.1016/S0010-0277(98)00058-4

Wu, Y. & Gweon, H. (2019). Surprisingly unsurprising! infants' looking time to improbable events is modulated by others' expressions of surprise. https://doi.org/10.31234/osf.io/8whuv

Wu, Y. & Schulz, L. (2018). Inferring beliefs and desires from emotional reactions to anticipated and observed events. *Child Development*, *89*(2), 649–662. https://doi.org/10.1111/cdev.12759

Wynn, K. (1992). Addition and subtraction by human infants. *Nature*, *358*(6389), 749–750.

Xu, F. (2019). Towards a rational constructivist theory of cognitive development. *Psychological Review*, *126*(6), 841. https://doi.org/10.1037/rev0000153

Xu, F. & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences*, *105*(13), 5012–5015. https://doi.org/10.1073/pnas.0704450105

Xu, F. & Griffiths, T. L. (2011). Probabilistic models of cognitive development: Towards a rational constructivist approach to the study of learning and development. https://doi.org/10.1016/j.cognition.2011.06.008

Xu, F. & Kushnir, T. (2013). Infants are rational constructivist learners. *Current Directions in Psychological Science*, *22*(1), 28–32. https://doi.org/10.1177/0963721412469396

Yogman, M., Garner, A., Hutchinson, J., Hirsh-Pasek, K., Golinkoff, R. M., Baum, R., Gambon, T., Lavin, A., Mattson, G., Wissow, L. et al. (2018). The power of play: A pediatric role in enhancing development in young children. *Pediatrics*, *142*(3). https://doi.org/10.1542/peds.2018-2058

Yoshikawa, H., Weiland, C., Brooks-Gunn, J., Burchinal, M. R., Espinosa, L. M., Gormley, W. T., Ludwig, J., Magnuson, K. A., Phillips, D. & Zaslow, M. J. (2013). Investing in our future: The evidence base on preschool education.

Yuniarto, L. S., Gerson, S. A. & Seed, A. M. (2020). Better all by myself: Gaining personal experience, not watching others, improves 3-year-olds' performance in a causal trap task. *Journal of Experimental Child Psychology*, *194*, 104792. https://doi.org/10.1016/j.jecp.2019.104792

Ziemann, L., Swaboda, N. & Ruggeri, A. (in preparation). Ecological learning in toddlers and preschoolers.