Goals, Play, and Cognitive Pragmatism: A study of flexible human minds

by

Junyi Chu

B.S., Vanderbilt University (2015)

Submitted to the Department of Brain and Cognitive Sciences
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY IN BRAIN AND COGNITIVE SCIENCES

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2023

Author . . . .
Department of Brain and Cognitive Sciences
MAY 19, 2023

Certified by . . . .
Laura E. Schulz
Professor of Cognitive Science
Thesis Supervisor

Accepted by . . . .
Mark T. Harnett
Graduate Officer, Department of Brain and Cognitive Sciences

Goals, Play, and Cognitive Pragmatism: A study of flexible human minds

by

Junyi Chu

Submitted to the Department of Brain and Cognitive Sciences
on MAY 19, 2023, in partial fulfillment of the
requirements for the degree of
DOCTOR OF PHILOSOPHY IN BRAIN AND COGNITIVE SCIENCES

Abstract

Few phenomena in childhood are as compelling or mystifying as play. While many animals play, human play is distinguished by the sheer diversity of goals that we pursue, even as adults. Yet the seeming inutility of play belies one of the hallmarks of intelligence: a remarkably flexible ability to reason and plan in novel situations. What kind of mind generates and pursues so many goals, and has so much fun in the process? In this dissertation, I suggest that answering this question requires us to go beyond current accounts of rational action and exploration. To map out the path forward I present three lines of research involving behavioral experiments with young children (ages four to six years) and adult comparisons. In study one I find that adults and children endorse speculative conjectures, even when implausible or lacking evidence, because we primarily evaluate novel proposals based on how well it answers our questions. In study two I demonstrate that children at play spontaneously take unnecessarily costly actions and pursue prima facie inefficient plans, even though they minimize costs when achieving similar goals in non-play contexts. Finally, study three demonstrates that adults and children value their goals from the moment they are chosen: participants stick with their goals even when less costly alternatives are available. On their own, each study contributes novel empirical findings and theoretical insights to their respective literature in explanation, play, and planning. Taken together however, they suggest a broader conclusion: that humans treat goals as valuable constraints for reasoning and decision-making. By paying attention to the goals we adopt and the problems we make for ourselves, we may explain much more of the richness and flexibility of the human mind.

Thesis Supervisor: Laura E. Schulz
Title: Professor of Cognitive Science

Acknowledgments

It is an immense privilege to complete this PhD with the support of many incredible people.

First thanks goes to my advisor, Laura Schulz. I cannot begin to count the times you saw a path forward when I was feeling lost among the trees. Your mentorship has enriched my understanding of science, of writing, and of life. It's been such a treat and honor to learn this craft and to think new thoughts alongside you. Thanks for always saying the right things, for opening doors, and for showing me the joy in science.

To my committee members, thank you for so many constructive critiques. Josh Tenenbaum, thank you for being around no matter how busy you were. I'm always impressed by how to have so many ideas and how you seem to have so much fun. I will always strive to emulate your unfailing memory, your respect towards ideas and people, and your ability to think so much, with so little, so quickly. Rebecca Saxe, thank you for demonstrating how to be robust, rigorous, and precise in both science and communication. I cherish our conversations about the human side of science. Caren Walker, thank you for always being an enthusiastic supporter and asking questions that help me learn by thinking. I look forward to continue conversations. I also want to thank Roger Levy for serving on my qualifying exam committee when some of my questions were still in their infancy, for writing my go-to reference article on mixed effects regressions, and for always taking students seriously.

Next, I want to thank my mentors at Vanderbilt and UCSD who prepared and encouraged me to pursue a PhD. Bethany Rittle-Johnson, thank you for two formative years studying how children learn math and walking me through every step of completing a research project. Georgene Troseth, thank you for the best introduction to cognitive development and writing through Harry Potter, for opening your lab to me as a wide-eyed sophomore, and for suggesting I write a thesis with Bethany. David Barner and the LADLab, thank you for teaching me that that simple ideas can be complex (red, yesterday, one, no), that complex ideas can become clearer (bootstrapping, implicatures, lme4), and for making me feel valued in all spaces.

I am so grateful to have "grown up in science" in 46-4011 with the Early Childhood Cognition Lab. In particular, I thank my lab buddies, Maddie Pelz for getting me started in years 1-4, and for bike rides to Harvard even in the rain, and Nicole Coates, for seeing me through the finish line in years 5-6. You two have seen every side of this research program, and me, through good times and harder ones. Thanks for being a sounding board for science and life and for all the coffee & matcha lattes. Thanks to Kim Scott for taking me on as a summer intern and for wise ways to think about statistics, code, and time. Thanks to Rachel Magid, Julia Leonard, and Yang Wu for making me feel welcome from day 1. Thanks to many past and present lab members for feedback, support, and numerous rounds of sing-a-longs and powerpoint karaoke: Max Siegel, Melissa Kline Struhl, Sophie Bridgers, Shengyi Wu, Izabele Jonusaite, Maya

Taliafero, Jessica Chomik-Morales, Kiera Parece, Pedro Tsividis, Julian Jara-Ettinger, Hyoweon Gweon, Elizabeth Bonawitz. Thanks to Sienna Radifera, Karolina Cabrera, Regina Ebo, Kary Richardson, and Sammy Floyd for making everything work.

For the work included in this thesis I am especially grateful to a group of dedicated research assistants who spent countless hours, sometimes on weekends, to work with participants, annotate data, and refine stimuli. I owe Chapter 3 to Rucha Kelkar, Cindy Zhou, Jaemarie Solyst, and Grace Cowles; Chapter 4 to Lucy Fu and Olivia Valle; and Chapter 5: Sofia Riskin, Bianca Santi, Ashley Lederman, Kameron Garland, and Asmita Mittal. Many thanks also to Sophia Diggs Galligan, Heyang Yin, Heidi Li, and Jinger Chong, and Christopher Montejo for their work on pilot and ongoing studies that did not make it into this dissertation, as well as the ever-growing Baby Physics team for 30fps precision annotations: Andre Weiss, Beyza Ciftci, Claire Ma, Cynthia Lei, Crystal Liu, Elisa Dimagiba, Faith Choe, Gabriel Kane, Grace Zhang, Hyunjin Christine Lee, Kailande Cassamajor, Katherine Johnson, Kevin Wen, Leensyn Rivera, Lillian Switkes, Liora Jones, Monica Correia, Naomi Kimiri, Salina Musyaju, Yuka Machino, Zoe Price. I learnt so much from each of you, and working together brought much joy and meaning to this work especially during the peak of the pandemic. I also owe many thanks to the Boston Children's Museum, Lookit, and ChildrenHelpingScience.com, organizations whose support and resources made possible data collection with hundreds of children over the years. Thank you for making this work possible.

My PhD has been nourished through interactions with a scientific community full of curious, motivated, and generous minds. I am indebted to: my cohort, for sticking around through it all; members of CoCoSci and SaxeLab for providing moral support of the computational and social kind; GPS, for working through what pragmatism means; Philosophy Circle, for surprising and thought-provoking conversations; Cambridge Writing Group, for kindness, accountability, and constructive feedback. Thanks to Jon Gauthier and Sam Cheyette for computational collaborations and making rough ideas more precise, and to Kelsey Allen, Peng Qian, Lionel Wong, Max Nye, Tan Zhi-Xuan, Desmond Ong, and Rico Rodriguez for many patient explanations about code and computation. Special mention to Desmond for finding my poster at my first CogSci in 2015 and for many great conversations over the years, to Alex Ferguson for writing momentum, and to Shari Liu, Ashley Thomas, Jenny Wang, and Pierina Cheung for honest and grounding conversations about careers. I'm grateful for my collaborators in science outreach and community building projects: Halie Olson and Heather Kosakowski for co-teaching Babies and Brains, Peng Qian and Karen Pang for co-organizing Connecting Minds, Carina Kauf for brainstorming Mind Games / BCS Challenges. Outside of MIT, thanks to CBMM 2018 and DISI 2022 for six beautiful summer weeks of amazing cross-disciplinary conversations, and thanks to the cogdev crew I look forward to seeing at every conference: Liz Lapidow, Rosie Aboody, Emily Liquin, Mika Asaba, Holly Huey, Minju Kim. More recently, to CoCoDev and CogToolsLab for 'adopting' me and helping me see the possibilities of this research program. All of you make science so much fun!

I want to thank my friends who make life outside the PhD just as rewarding: Martin and Chilla, Jungsoo, Jon, Zidong, QingQing and Gerver, Chen Jia, Hong and Gawa, for memories, honesty, and joy. Thanks to my best friend, Yiran for staying connected throughout the decades. You are my rock. Thanks to my partner, Renz, and our dog Whiskey, for inexhaustible support and unconditional love. Thanks to Alex for always playing along, and for making me proud to be a big sister. Thanks to my aunt Weiping and uncle Chun for a home away from home, for Thanksgiving meals and trips to National Parks. I love you all very much.

Most importantly: Thanks to my parents, Jane and Lixin, for teaching me values of learning, joy, and kindness, and for always being there for me from day 0, no matter how far apart we are.

This dissertation includes material that has been previously published. I am grateful to reviewers and editors for their helpful comments.

- Chapter 2 is a reprint of the material as it appears in Chu, J. & Schulz, L. E. (2020) Play, curiosity, and cognition. *Annual Reviews of Developmental Psychology*, 2:317-343.

- Chapter 3 is a reprint of the material as it appears in Chu, J. & Schulz, L. E. (2021) Children selective endorse speculative conjectures. *Child Development, 92*, 6, e1342-e1360. Permission for reuse of material has been obtained from Wiley.

- Chapter 4 is a reprint of the material currently in press as Chu, J. & Schulz, L. E. (in press) Not playing by the rules: Exploratory play, rational action, and efficient search. *Open Mind*.

- Chapter 5 includes material previously published as Chu, J. & Schulz, L. E. (2022) "Because I want to": Valuing goals for their own sake. In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*. 1263-1269.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

> I do not know what I may appear to the world, but to myself I seem to have been only like a boy playing on the sea-shore, and diverting myself in now and then finding a smoother pebble or a prettier shell than ordinary, whilst the great ocean of truth lay all undiscovered before me.
>
> Sir Isaac Newton

This dissertation considers the flexibility of human goal pursuit, which you can see in action at any playground. Take a look - do you see a slide, and children taking turns to go down it? What about someone trying to run up the slide, or go down backwards? Stay a little, and walk around. You might catch a three-year-old having the time of their lives playing with a stick. Try to predict what they might do next: are they drawing in the sand box, having a sword fight, or "learning how to turn things into pigs" (Colliver & Fleer, 2016)? There appears to be an unbounded set of ways that children amuse themselves. How do children come up with these activities, and how do these activities keep children so engrossed?

Such behaviors don't disappear in adulthood. Try asking someone how they spend their free time: Maybe they have a hidden talent for flipping a water bottle and landing it right side up no matter how full it is. Or you might meet someone who has cycled across the country, who performs stand up comedy about science, or who is fluent in the fictional language of Na'vi. A quick browse through your favorite social media platform or the Guinness book of world records will reveal many more examples of dedication to seemingly arbitrary goals. The demands of modern living may give some common structure to how we spend our time and energies, but left to our own devices, we may derive a significant amount of value in the pursuit of idiosyncratic goals and personal passions.

What kind of mind yields the rich diversity of goals that humans entertain and adopt, as well as the flexibility with which we reason and plan in pursuit of these goals – not to mention the fun we get out of doing so? This is the central question underlying my research program. By studying the playfulness of human minds, I believe we can better understand the nature of human cognition, innovation, and motivation. I suspect a full answer to this question would require a long-term research program in cognitive science. This program ought to generate satisfying and mutually compatible explanations for each of these four major questions (c.f. Tinbergen, 1963):

1. **Function: What is it for?** Having a flexible reasoning and planning system is arguably a core part of human intelligence in that this enables us to solve a wide range of problems, and adapt to new situations. But it seems harder to explain the function of goals which may be arbitrary and unlikely to pay off with respect to immediate benefits to survival or any predictably useful learning. What benefit to fitness does the pursuit of arbitrary goals confer? Or, in the spirit of rational analysis (Anderson, 1991; Chater and Oaksford, 1999), what problem is being solved by a cognitive system that generates its own goals and problems?

2. **Mechanism: How does it work?** What computations and representations are involved in generating and selecting new goals and plans? How are these implemented in the brain?

3. **Ontogeny: How did it develop within individuals?** How does the capacity for flexible goal generation and pursuit develop over the lifespan? Are children more or less creative, flexible, and dedicated to these arbitrary goals than adults? What role does learning have to play in how we generate, choose, and solve novel problems? Are the motivating factors driving the arbitrary projects in adulthood the same factors behind children's propensity play?

4. **Phylogeny: How did it develop over evolutionary history?** Many non-

human animals play, and many species display behavioral flexibility in novel situations (e.g., some birds are well known for solving novel physical reasoning problems). What aspects of human goal-directed behavior is shared and distinct from other species? What selection pressures gave rise to these features?

These are big questions. In this dissertation I take a first step, by laying an initial groundwork of empirical and theoretical observations.

From rational to playful minds

Research in the psychological and cognitive sciences have made immense progress towards explaining how humans reason and act. One dominant approach considers human intelligence as the product of a *rational* mind: one that can assess various goals and choose actions that maximize expected utilities, i.e., aiming to achieve the highest rewards with lowest costs (Anderson, 1991; Chater & Oaksford, 1999; Gershman, Horvitz, & Tenenbaum, 2015; Lieder & Griffiths, 2020). This principle applies in both physical contexts (e.g., preferring one cookie on a nearby plate over one in a closed container on a high shelf; Jara-Ettinger, Gweon, Schulz, & Tenenbaum, 2016; Jara-Ettinger, Gweon, Tenenbaum, & Schulz, 2015) and in social interactions (e.g., preferring concise speech, or efficiently distributing tasks across agents; Frank & Goodman, 2012; Magid, DePascale, & Schulz, 2018; Mascaro & Csibra, 2022)

In addition to acting in rational ways, we also make rational inferences: we update our beliefs by weighing the evidence from observed data against our prior knowledge (Gopnik & Wellman, 2012; Tenenbaum, Griffiths, & Kemp, 2006). Combine this with a rich set of initial representations (Spelke, 2022), powerful learning mechanisms that attend to both observable statistical regularities and unobservable causal and abstract relations (e.g., Carey & Spelke, 1994; Saffran, Aslin, & Newport, 1996; Schulz, 2012b), plus an intrinsic drive to gain information (Berlyne, 1966; FitzGibbon, Lau, & Murayama, 2020; Gottlieb & Oudeyer, 2018; Kidd & Hayden, 2015; Loewenstein, 1994), and we get a mind built to acquire accurate models of the external world – far more quickly, efficiently, and robustly than we can currently engineer.

Notably, human minds are not perfect; our behavior often deviates systematically from the predictions of these rational models (e.g., Tversky & Kahneman, 1974). Research in economics, psychology, neuroscience, and computational cognitive science have been able to explain many of these deviations by recognizing the important ways that our minds are constrained (Hilbert, 2012; Simon, 1955), subject to limitations of memory, time, and other computational resources (Gershman et al., 2015; Lieder & Griffiths, 2020; Schacter, 2001), as well as adapted to the physical and social environments we live in (Anderson, 1990; Gigerenzer & Brighton, 2009; Todd & Gigerenzer, 2012).

However, I suggest there are still some important gaps in a full explanation of human behavior that accounts of rational learning and rational action leave open – gaps which we find when paying attention to the playful behaviors at the beginning of this chapter. In this thesis, I explore three such gaps, grounding each in a different empirical case study.

First, a mind focused on epistemic goals may get the world right, but what if available facts do not answer our questions or solve our problems? While many lines of inquiry in science began with a surprising observation, or by exploring the potential application of new tools, much of scientific progress has also come from posing questions ahead of having any pertinent data or method in hand. In these situations, how could one assess which questions and conjectures to are worth entertaining? We explore this question in **Chapter 3**, looking at how children and adults evaluate novel conjectures to questions that cannot be resolved with available data.

Second, how might rational cost-benefit analyses account for the more arbitrary and playful behaviors observed in humans? In **Chapter 2**, I review cognitive and non-cognitive accounts of play, and point out what appears to be an incompatibility between existing accounts and the full richness of children's play. On the one hand, lab studies of children's exploratory play have shown that their spontaneous behaviors reflect rational processes of learning about the world: Children are sensitive to opportunities to learn, and in play, can spontaneously learn about causal systems and abstract relations governing the system that they are exploring. On the other hand, real-world observations of children playing are often harder to explain as maximizing the same set of information and practical utilities. I suggest that, in addition to studying play as rational exploration for epistemic ends, a distinctive characteristic of play is the generation of new goals and the creation of new problems, which may be pursued regardless of expected learning gains. I test this proposal in **Chapter 4**, with three experiments showing that children's exploratory play is indeed distinguished from functional behavior by the prevalence of seeming violations of rational, efficient action.

Third, not only do we make up new problems and solutions, but we also value them, even if the problem is not feasible or the solution not practical. How can we explain our attachment to and continued pursuit of goals which may not pay off in the short term or may even be impossible to achieve within a lifetime? One example may suffice. Historians of the Tang Dynasty note that at least six emperors died from taking (presumably prototype) "elixirs" of immortality (Chiang, 2007). Surely the first few accidents would have deterred future attempts. Yet the pursuit of immortality continued, both in alchemical experimentation (which eventually led to the synthesis of gunpowder) and in the development of Taoist spiritual practices. **In Chapter 5**, I develop a simple paradigm (in a much safer context) asking whether even young children stick to their goals when alternatives are less costly and just as valuable.

More generally, while current accounts of rational learning and rational action rely on the dual engines of data (observations of the world) and theory (our prior beliefs and conceptual structures), I propose that we must also appeal to a third constraint: *goals*. This approach is inspired by research in cognitive science that considers how non-epistemic goals can help explain more of the complexity of human behavior. In communication, for example when someone asks for feedback, we may trade off between being polite and being accurate (Yoon, Tessler, Goodman, & Frank, 2020). In moral decisions, we may balance loyalty and fairness (Waytz, Dungan, & Young, 2013). And in exploration, our goals shape what we pay attention to. We may look at different parts of a scene depending on what question we are hoping to answer about it (Yarbus, 1967) or what we think a co-observer is talking about (Cooper, 1974; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995).

Thus we arrive at the main thesis in this dissertation: **that by treating the goals we want to achieve and the problems we want to solve as additional computational constraints on cognition, we might be better able to explain the richness and flexibility of human thought.**

This particular approach is also inspired by the philosophical tradition of pragmatism in epistemology (Dewey, 1916; James, 1907; Peirce, 1878) which attempted to introduce a new way of thinking about knowledge beyond rational first principles and empiricist observations. While acknowledging that we are certainly attentive to notions of objective truth and probability, pragmatists also recognized the value of ideas in terms of how expedient it is. That is to say, we may value ideas which are useful with respect to our goals.

However, beyond claiming that goals *affect* our reasoning and attention, I am also interested in exploring two stronger claims: (1) that goals are a *central* organizing element to human cognition, without which we might face fundamentally intractable problems of search and decision-making, and (2) goals are sources of value in their own right.

Plan for the dissertation

In this dissertation I investigate the thesis that humans treat goals as critical computational constraints on inference and action. In the chapters to follow, I present research spanning play, explanation, and decision-making, to explore the different ways that human cognition is organized around and motivated by the goals we adopt.

I begin by reviewing the literature on play, curiosity, and cognition in both human and non-human animals in **Chapter 2**, with special attention to accounts of the relationship between play and cognition in early childhood. While much of children's exploratory play can be aligned with accounts of rational exploration, thus explaining the richness and flexibility of children's *learning*, I argue that these accounts fall short of explaining full richness and flexibility of *play*, especially the kinds of play that emerge after early childhood and which persist into adulthood. Instead, I suggest that a distinctive characteristic of human play is our ability to set flexible goals and make up new problems (*play for problems and proposals*), and discuss how this kind of goal-directed exploration contrasts with existing accounts of play as involving rational exploration and information-seeking curiosity.

Next, in **Chapters 3 to 5** I present three empirical case studies each exploring different features of humans' ability to pursue and create goals that go beyond existing accounts of rational learning and rational action. All three chapters follow a common approach: in each study, we design tasks where existing accounts of rational choice make a clear prediction about which options participants should prefer, such as evaluating explanations, planning to achieve a target outcome, or making choices between two options. However, by experimentally manipulating participants' goals, we find that responses systematically shift away from these predictions – which rely on environmentally defined measures of information gain, action costs, and expected outcomes – and instead, shift towards what I will term "alignment with goals". In addition, because we are interested in aspects of human cognition that may explain arbitrary goal pursuits in both adults and young children, all three chapters include children between about three to seven years old as participants. We find mostly similarities between children and adults in these initial studies. Together, these studies provide evidence for rich, flexible, and ad hoc reasoning and planning abilities in both adults and young children.

In **Chapter 6**, I build off the theoretical review and empirical findings in previous chapters to develop a proposal about the function of flexible goal pursuit. I suggest that humans are intrinsically motivated to seek not just opportunities to gain information or achieve outcomes, but also, opportunities to engage in the activities of reasoning and planning. Thus, flexible human goals, independent of the value of expected outcomes, may attract our cognitive engagement insofar as they sustain (or we expect them to sustain) thinking.

Finally, **Chapter 7** concludes with a discussion of the implications of this thesis. What if goals aren't just the ends that we reason about and plan for, but instead, are the critical structures that organize our mental models and processes and without which, thinking would be very hard? What if the goals we adopt don't just reflect our values, but are themselves sources of value that lead us to engage in decision-making and planning in the first place, and to begin thinking at all? I discuss how the view of goal-centered thinking developed in this thesis might help us better understand other phenomena in human cognition.

We begin with asking what children's play can teach us about human cognition.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

Play, curiosity, and cognition

> It is a happy talent to know how to play.
>
> — Ralph Waldo Emerson

Abstract

Few phenomena in childhood are as compelling – and mystifying – as play. Here we review five proposals about the relationship between play and development. We believe each captures important aspects of play across species; however, we believe that none of them accounts for the extraordinary richness of human play, or its connection to distinctively human learning. In thinking about play, we are particularly struck by the profligacy with which children set seemingly arbitrary rewards and incur unnecessary costs. We suggest that researchers take the seeming inutility of play seriously, and consider why it might be useful to engage in "useless" behavior. We propose that humans' ability to choose arbitrary costs and rewards allows us to pursue novel goals, discover unexpected information, and invent problems we wouldn't otherwise encounter. Because problems impose constraints on search, these invented problems may help solve a big problem: the problem of how to generate new ideas and plans in an otherwise infinite search space.

2.1 Introduction

Play is one of the most enchanting and baffling phenomena in nature. Among the most accessible of all behaviors, it is also among the most difficult to characterize rigorously. We all recognize play when we see it. Nonetheless, play eludes definition to the extent that a species of play (games) has served to illustrate the limits of classic theories of word meaning (Wittgenstein, 1953/2001).

However, if play has defied description, it is not due to lack of study. Scientists in fields ranging from ethology to robotics have debated the factors that might motivate play and the functions that play might serve. In this article, we review both noncognitive and cognitive accounts of play, focusing especially on recent research linking play, epistemic curiosity, and learning. Reflecting the traditions in which these accounts are best developed, we review research mostly on nonhuman animals in discussing noncognitive accounts and mostly on humans in discussing cognitive accounts, but the accounts apply across species and are not mutually exclusive: Play might emerge for many reasons, serve many ends, and occur in different forms in a single play session.

Ultimately, however, we conclude that none of the current accounts does justice to the richness of distinctively human play – or distinctively human curiosity and cognition. We argue that understanding play in human beings requires taking its apparent uselessness seriously. Indeed, we suggest that among the most salient features of human play is the degree to which we intervene on our own utility functions. That is, in play, humans willingly adopt arbitrary rewards and incur unnecessary costs, leading to systematically different behavior in play than in other forms of intentional, goal-directed behavior. We suggest that this willingness to incur unnecessary costs to achieve idiosyncratic ends allows humans to create a vast array of problems we would not otherwise have. We propose that these invented problems, and the constraints they impose, help solve a big problem: how to generate new ideas and plans in an otherwise infinite search space.

2.2 Non-cognitive accounts of play

2.2.1 Play for pleasure

We begin by discussing noncognitive accounts of play, starting with what is surely the simplest possibility: that play has no function at all. We may play just for the pleasure of it. More precisely, insofar as animals evolved to find behaviors that increase reproductive success rewarding, they may be expected to engage in these behaviors often, even in contexts where they confer no advantage. Thus, if it is rewarding for dolphins to blow bubbles into nets to hunt fish (Ingebrigtsen, 1929; Jurasz & Jurasz, 1979; F. A. Sharpe & Dill, 1997), they may also blow bubble rings

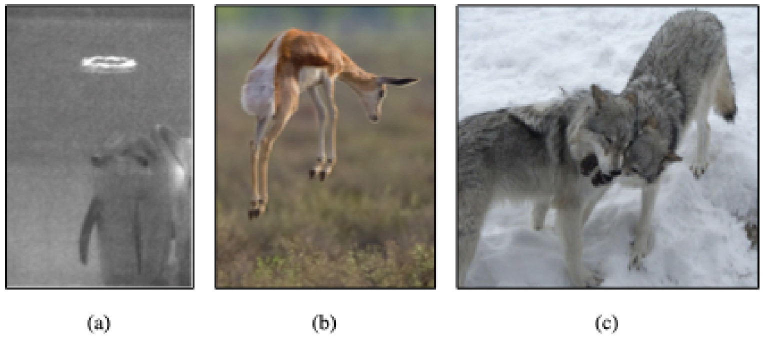

Figure 2-1: Examples of the kinds of behaviors associated with each of the non-cognitive accounts of play (a) Play for pleasure: A dolphin blowing bubble rings (McCowan et al., 2000); (b) Play for performance: A springbok showing off its youth and fitness by pronking (bouncing off all 4 legs); (c) Play for peace-making: Play fighting in wolf cubs possibly as a low cost way to establish dominance hierarchies; photo by Zechariah Judy.

just for fun (**Figure 2-1a**; see Delfour & Aulagnier, 1997; McCowan et al., 2000; Pace, 2000); if it is positively arousing for chimpanzees to swing, jump and leap to travel through the forest, they may do so just for the pleasure of it (Mears & Harlow, 1975); and if a preference for colorful, soft substances allowed primates to detect ripe fruit (e.g., Dominy, Garber, Bicca-Marques, & Azevedo-Lopes, 2003), they may then enjoy playing with colorful, squishy things in any context (witness the 280 million entries on Google associated with the current “slime” craze).

These are of course “just so” stories; here however, they serve as “just not so” stories, explaining not why an observed behavior fulfills an adaptive end but why it may not – why playful behaviors may simply be generalizations of behaviors that are functional in other contexts. In this sense, play may be an evolutionary spandrel (Gould & Lewontin, 1979), persisting only because it is reinforced by reward systems evolved for other purposes. Of course, our difficulty imagining how some behaviors – blowing bubble rings, or playing with slime – could be useful does not mean that no such use exists. The anthropologist Robin Dunbar cautioned against the "Spandrel Fallacy": 'I haven't really had time to determine empirically whether or not something has a function so I'll conclude that it can't possibly have one.' (Dunbar, 2012). Still, it is hard to know what evidence could disconfirm the possibility that, at least in some contexts, animals "play for pleasure".

2.2.2 Play as performance

Animals may play not (only) because play is rewarding but also, paradoxically, because play is costly, both in terms of time and energy and in terms of risks to life and limb (Harcourt, 1991; L. L. Sharpe, Clutton-Brock, Brotherton, Cameron, & Cherry, 2002). Animals (including human children) play only when they are healthy, well fed, and safe, and they stop playing when they are injured or under stress (Alessandri, 1991; Burghardt, 2005; Dawkins, 2006; R. M. Fagan, 1981; Fagot & Kavanagh, 1991; Fraser & Duncan, 1998; Held & Špinka, 2011; Lawrence, 1987; Martin & Caro, 1985; Spinka, Newberry, & Bekoff, 2001) [1]. Since play is both costly and easy to observe, it may function as an honest signal of health and fitness (e.g. **figure 2-1b**).

Moreover, play is also a sensitive signal: it drops off quickly in response to real and perceived threats, recovers quickly in their absence, and flourishes in resource-rich environments (R. M. Fagan, 1981; Panksepp & Burgdorf, 2010; P. Thornton & Waterman-Pearson, 2002). For instance, baboons' play closely tracks annual rainfall (Barrett, Dunbar, & Dunbar, 1992) and meerkats' play doubles relative to controls when their food is supplemented (L. L. Sharpe et al., 2002). Insofar as predators may be less likely to attack (and conspecifics more likely to mate) with animals who look like they are uninjured, well-fed, and vigorous, play might be favored by both natural and sexual selection [2]. We will refer to the idea that play might function as a signal of fitness as "play for performance".

2.2.3 Play for peace-making

Play might enhance fitness, not simply advertise it. In particular, researchers have suggested that social play might reduce within-group aggression and increase within-group coordination. Pack and herd animals who can evaluate one another's strength and establish dominance hierarchies through play might be more likely to avoid riskier

[1] There are some exceptions however, where increased stress leads to increased play. Kittens and rat pups weaned earlier than usual, and yearling rhesus monkeys deprived of care by the birth of a sibling, play more than their age-mates, arguably as a step towards increased independence (Bateson, Martin, and Young 1981; Bateson, Mendl, and Feaver 1990; Devinney, Berman, and Rasmussen 2003; E. F. Smith 1991, see Held and Špinka 2011 for discussion)

[2] Since play is associated with juveniles, it might be difficult to imagine a role for it in mate selection but in fact play persists into adulthood in humans cross-culturally (Roberts & Sutton-Smith, 1962), and in most other animals observed, including rhesus monkeys (Breuggeman, 1978); horses (Hausberger, Fureix, Bourjade, Wessel-Robert, & Richard-Yris, 2012), cats (S. L. Hall & Bradshaw, 1998), cormorants and herons (Sazima, 2008); otters (Beckel, 1991), bottlenose dolphins (Kuczaj & Eskelinen, 2014) and humpback whales (indeed, the latter two have been observed playing with each other, Deakos et al. 2010.)

fights that could weaken the group as a whole (Dolhinow, 1999; Palagi, 2006, 2008; Panksepp, 1981; Pellis & Iwaniuk, 2000; Pellis & Pellis, 1991; P. K. Smith, 1982; Thompson, 1998; Zimen, 1982). Additionally, attention to the metacommunicative signals used in play (exaggerated calls and postures, repeated movements like head-shaking and tail-wagging; Bekoff, 1972) might support social attunement and greater cooperation in hunting prey and fending off predators (**Figure 2-1c**).

However, although the idea that "animals that play together stay together" (Bekoff, 1974) has been influential (Baldwin & Baldwin, 1974; Bekoff & Byers, 1985; Berman, 1982; Drea, Hawk, & Glickman, 1996; Gaines & McClenaghan Jr, 1980; K. Hall, 1968; Holmes, 1995; Jay, 1963; P. C. Lee, 1982; Panksepp, 1981; Poirier, 1969; Poirier & Smith, 1974), evidence for the claim is mixed. Some research has found that play (e.g., in coyote pups) is inversely correlated with sibling aggression (Drea et al., 1996) but many other studies that have looked for relationships between play and positive social outcomes have failed to find them. Thus, for instance, individual differences in juvenile play have no effect on within-group aggression or social dispersion in wallabies (Watson, 1993); squirrel monkeys (Baldwin & Baldwin, 1974); wolves (Cordoni, 2009); rats (Pellis & Iwaniuk, 1999); or meerkats (L. L. Sharpe, 2005; L. L. Sharpe & Cherry, 2003). Arguably however, play might still have species-level effects on peace-making insofar as highly intelligent social species often display high levels of social cohesion despite also having high levels of within-species aggression (De Waal, 1986).

2.3 Cognitive accounts of play

2.3.1 Play for practice

Figure 2-2: Play for practice. Rock juggling in otter pups possibly fosters the motor skills needed to use rocks to open mollusks as an adult (though evidence that this play really does support adult skills has been hard to come by; see Allison et al., 2020).

Thus far we have discussed non-cognitive accounts of play; we now turn to the idea that play supports learning. In all of modern psychology, perhaps few claims are so uncontroversial – and so hard to substantiate. Parents, educators, and researchers alike believe that play in early childhood supports learning (Berlyne, 1969a; Bruner, Jolly, & Sylva, 1976; Golinkoff, Hirsh-Pasek, et al., 2006; Groos, 1901; Piaget, 1962; Vygotsky, 1934/1962), and across species, it is clear that smarter, more behaviorally flexible species play more (Bjorklund, 1997; Groos, 1898; Pellegrini, Dupuis, & Smith, 2007). Nonetheless, establishing specific relationships between play and learning remains a challenge.

The most straightforward way that play could support adult behavior is not through learning but by increasing physical fitness (Bekoff, 1988; Byers, 1998; R. M. Fagan, 1981). However, because exercise induced effects of fitness are transitory (see e.g., Byers, 1998), an alternative possibility is that play helps juveniles to master locomotor skills critical to adulthood (play as practice; 1898; see also Bekoff and Byers 1998; Burghardt 2005; R. M. Fagan 1981; Pellegrini et al. 2007). This seems especially plausible with respect to the complex motor skills involved in hunting or using tools. Thus, kittens might pounce on strings, chimps and crows play with sticks, and otter pups play with rocks in order to, respectively, be better able to catch mice, extract ants and larva from crevices, and crack mollusk shells as adults (**Figure 2-2**) (Caro, 1995; Humle, 2006; Inoue-Nakamura & Matsuzawa, 1997; Nishida & Hiraiwa, 1982; Rutz et al., 2010). Consistent with the idea that play is preparation for adult behavior, children, cross-culturally, are given scaled-down, often non-functional versions of adult tools as playthings (e.g. Gusinde, 1931; Healey, 1990; MacDonald, 2007; Watanabe, 1975), and babies spend many hours manipulating objects before they master the use of even simple tools like spoons or rakes (Connolly & Dalgleish, 1989; Lockman, 2000; McCarty, Clifton, & Collard, 1999, 2001; Piaget, 1952).

However, although the idea of play as practice for adult life is intuitive, there is surprisingly little evidence that play in juveniles correlates with skill in adults. Thus for instance, kittens raised without toys grow up to hunt as well as kittens surrounded with them (Caro, 1980), the amount of play meerkats engage in as youngsters is uncorrelated with their success in hunting or in territorial disputes as adults (L. L. Sharpe, 2005), and otters who juggle rocks more frequently are not any faster at extracting food (Allison et al., 2020). Similarly although there is considerable evidence that developmental delays and disorders affect exploratory play in humans (de Almeida Soares, von Hofsten, & Tudella, 2012; de Campos, da Costa, Savelsbergh,

& Rocha, 2013; Kaur, Srinivasan, & Bhat, 2015; Kavšek, 2004; Kavšek & Bornstein, 2010; Kopp & Vaughn, 1982; Koterba, Leezenbaum, & Iverson, 2014; Loveland, 1987; Ruff, McCarton, Kurtzberg, & Vaughan Jr, 1984; Sigman, 1976; K. P. Wilson et al., 2017; Zuccarini et al., 2016), there is only weak evidence that typical exploratory behavior correlates with later outcomes (Bornstein, Hahn, & Suwalsky, 2013; McCall & Carriger, 1993; Muentener, Herrig, & Schulz, 2018; Raine, Reynolds, Venables, & Mednick, 2002; Viholainen et al., 2006). Moreover, the further removed the juvenile behavior is from motor coordination, the less compelling the relationship between play and adult skills becomes. A recent meta-analysis for instance found no strong evidence for causal relationships between pretend play and cognitive outcomes for any of the areas (intelligence, creativity, problem solving, theory of mind, language, executive function, and emotion regulation) for which links had been proposed (Lillard, 2012). Nonetheless, some skills, especially those related to implicit skill learning (i.e, playing instruments, playing sports) are clearly easier to learn before late adolescence than later in life (Janacsek, Fiser, & Nemeth, 2012). Thus for at least some kind of behaviors, the particular activities practiced in juvenile play are indeed likely to have enduring impacts.

2.3.2 Play for prediction and plans

The idea of play as practice suggests that most of the benefits of play are incurred in adulthood. However, the most influential current accounts of the relationship between play and learning suggest that play behaviors are motivated by learners' moment-to-moment epistemic curiosity. Further, these accounts suggest that information gained in play has online effects in reducing learners' uncertainty and in increasing their ability to predict events in the world. There have been several excellent discussions of epistemic curiosity, exploration and self-directed learning in the past decade (Gottlieb, Oudeyer, Lopes, & Baranes, 2013; Gureckis & Markant, 2012; Jirout & Klahr, 2012; Kidd & Hayden, 2015; Silvia, 2012) so we will not attempt another comprehensive review here. Instead we will focus on a few key findings in the developmental literature and highlight relevant connections to work in ethology, artificial intelligence (AI), robotics, and computational cognitive science.

In the developmental literature, the link between play and learning has largely focused on connections between exploratory behavior and children's causal reasoning. The earliest form of exploration we can measure is visual exploration, and looking time methods have yielded rich accounts of infant perceptual and cognitive abilities (for review, see Aslin, 2007; Haith, 1980). Like many other animals, infants preferentially look at stimuli that are novel (J. F. Fagan, 1970; Fantz, 1964; Saayman, Ames, & Moffett, 1964), perceptually salient (Civan, Teller, & Palmer, 2005; Kaldy & Blaser, 2013), and relatively complex (Brennan, Ames, & Moore, 1966; Cohen, 1972; Thomas, 1965) and this selective attention may help infants learn statistical properties across events ranging from patterns of shapes to phonetic alternations (Fiser & Aslin, 2002; Kirkham, Slemmer, & Johnson, 2002; Maye, Werker, & Gerken, 2002; Saffran et al., 1996; Saffran, Johnson, Aslin, & Newport, 1999; Saylor, Baldwin, Baird, & LaBounty, 2007).

Infant's visual attention is not just stimulus-driven; infants look longer at events that violate their expectations of the world (for reviews, see Spelke, 1985; Spelke, Breinlinger, Macomber, & Jacobson, 1992). Researchers using violation-of-expectation and preferential looking paradigms have made fundamental discoveries about infants' early representations of objects and forces (Baillargeon, 1987; Baillargeon, Spelke, & Wasserman, 1985; I. K. Kim & Spelke, 1992, 1999), number (McCrink & Wynn, 2004; Wynn, 1992), probability and sampling (Xu & Garcia, 2008), agents and goals (Liu, Ullman, Tenenbaum, & Spelke, 2017; Luo, 2011; Onishi & Baillargeon, 2005; Woodward, 1998), social interactions (Hamlin, Wynn, & Bloom, 2010; Powell & Spelke, 2013), and emotions (Skerry & Spelke, 2014; Y. Wu & Gweon, 2019; Y. Wu & Schulz, 2018).

Of course, visual exploration does not constitute play, per se. But by six months of age, infants begin to manually explore objects and the physical environment. While earlier researchers thought that infant exploratory play might be repetitive and perseverative (e.g. Piaget, 1954), recent work suggests that infants selectively explore objects that appear to violate their naive theories, and explore in ways specific to the apparent violation (for review, see Stahl & Feigenson, 2018). For instance, infants tend to drop toys that appear to violate gravity but bang toys that appear to violate solidity (Stahl & Feigenson, 2015). Moreover, infants do not engage in this kind of exploration if the apparent violation can be explained away (e.g., a toy appears to pass through a solid wall but the wall is then turned to reveal a gap; Perez and Feigenson 2022).

Children's exploration becomes increasingly sophisticated after infancy and throughout the preschool years (Pelz & Kidd, 2020). Toddlers can use co-variation evidence to determine the probable cause of failed actions and seek help or explore accordingly

(Gweon & Schulz, 2011), and will selectively explore objects depending on whether evidence for the extension of object properties was drawn randomly or selectively (Gweon, Tenenbaum, & Schulz, 2010). Children selectively explore evidence that violates their prior expectations (Bonawitz, van Schijndel, Friel, & Schulz, 2012; Schulz, Standing, & Bonawitz, 2008) and both toddlers and preschoolers will selectively search for unobserved causes given theory-violating evidence (Muentener & Schulz, 2014; Schulz, Hooppell, & Jenkins, 2008; Schulz & Sommerville, 2006; D. M. Sobel, Yoachim, Gopnik, Meltzoff, & Blumenthal, 2007). Preschoolers also explore and engage in active hypothesis testing given ambiguous or confounded evidence (Cook, Goodman, & Schulz, 2011; Schulz & Bonawitz, 2007; van Schijndel, Visser, van Bers, & Raijmakers, 2015) and recent work suggests that children's exploratory play is closely calibrated to their uncertainty, quantitatively varying with the difficulty of discrimination problems (Siegel, Magid, Pelz, Tenenbaum, and Schulz 2021; see **Figure 2-3**).

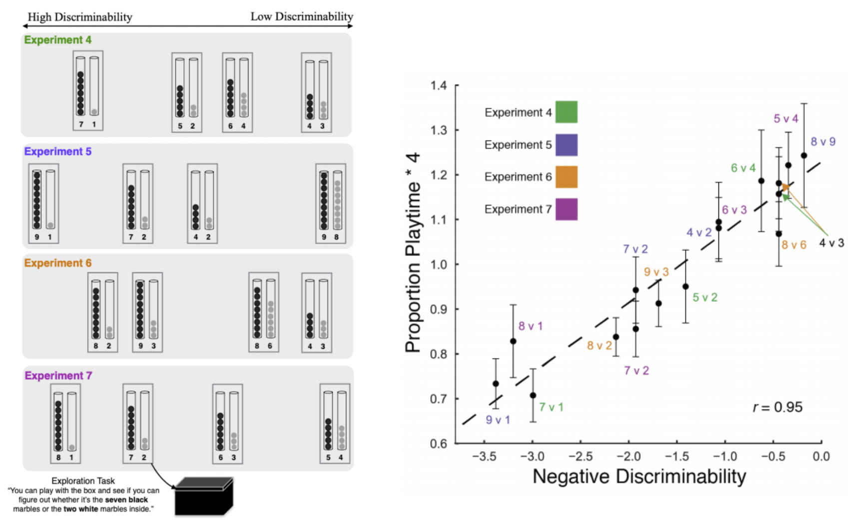

(a) (b)

Figure 2-3: Play for prediction and plans. (a) Children shook a box to guess how many marbles were inside; (b) Their exploration time tracked the difficulty of discriminating between the heard and unheard alternative in a remarkably fine grained way. Figure adapted from (Siegel et al., 2021)

And children's exploratory play supports causal learning (McCormack, Frosch, Patrick, & Lagnado, 2015; Schulz, Gopnik, & Glymour, 2007); even two and three-year-olds can discover abstract relations, including hierarchical causal structures, in free play (Sim, Mahal, & Xu, 2017). Children attend more to the effects of their own interventions than observed evidence (Fireman, Kose, & Solomon, 2003; Kushnir & Gopnik, 2005; Kushnir, Wellman, & Gelman, 2009), and in some cases, children may learn better through free play than through observation alone (D. M. Sobel & Sommerville, 2010). Moreover, the link between play and causal reasoning is not limited to exploratory play; some work suggests that children's pretend play also supports causal and counterfactual reasoning (Buchsbaum, Bridgers, Skolnick Weisberg, and Gopnik 2012, see also Gopnik and Walker 2013; Kavanaugh and Harris 1999; Weisberg 2015).

finally, children integrate causal and social information in their play. Both preschoolers' and toddlers' exploratory play is sensitive to whether evidence is provided accidentally or pedagogically (Bonawitz et al., 2011; Butler & Markman, 2012, 2014; Jean, Daubert, Yu, Shafto, & Bonawitz, 2019; Shneidman, Gweon, Schulz, & Woodward, 2016) and also whether evidence is selectively withheld (Gweon, Pelton, Konopka, & Schulz, 2014). Preschoolers explore more when adults provide information about the function of toys in the form of questions rather than statements (Yu, Landrum, Bonawitz, & Shafto, 2018), and selectively use the more informative of prior knowledge or social cues to guide their exploration (Luchkina, Sommerville, & Sobel, 2018). Children also use the results of their own exploration to teach others (Gweon & Schulz, 2018), and guided play by teachers supports children's learning while increasing their engagement (Bustamante et al., 2020; Fisher, Hirsh-Pasek, Newcombe, & Golinkoff, 2013; Weisberg, Hirsh-Pasek, & Golinkoff, 2013). Thus, although children's play might often appear random or haphazard, collectively this work suggests that children's play is connected to principles that could support learning and discovery in early childhood (for review, see Schulz 2012b).

We have restricted our review to the literature on children. Doing justice to the work on exploratory behavior and learning in non-human animals and in artificial agents is beyond the scope of this chapter. However, we want to highlight three common threads that unite discussions of play and learning across developmental, ethological, and computational approaches.

first, the motivation to seek new information is widespread. Humans do it – so do crows and chimps, octopi and orangutans[3] (Mather & Anderson, 1999; Welker, 1956; Wimpenny, Weir, & Kacelnik, 2010). Even the search behavior of animals as simple as moths and roundworms can be characterized by models of maximally informative foraging (Calhoun, Chalasani, & Sharpee, 2014; Vergassola, Villermaux, & Shraiman,

[3]Although interestingly, wild orangutans are far more neophobic than neophilic; only in captivity do orangutans show high rates of exploration.

2007). Moreover, the motivation to explore is robust. Animals will forego immediate, tangible rewards and incur costs, including physical pain, to gain information (e.g., hungry and thirsty rats will delay eating and drinking to explore new terrain, and rats conditioned to fear an electrified grille will cross it to explore; Nissen 1930; Zimbardo and Montgomery 1957). Some of this behavior can be characterized as instrumental behavior in which foregoing immediate rewards enhances overall gains in the longer term. But there is also a wealth of evidence that animals value information in its own right, even when it serves no instrumental end (Bennett, Bode, Brydevall, Warren, and Murawski 2016; Blanchard, Hayden, and Bromberg-Martin 2015; Gottlieb and Oudeyer 2018; Vasconcelos, Monteiro, and Kacelnik 2015, see also Pellegrini et al. 2007; Spinka et al. 2001).

Second, while the motivation to explore is early-emerging, widespread, and robust, it is not indiscriminate. Learners do not attend merely to the degree to which information is novel or unpredictable; if they did, they would spend much of their time exploring stimuli that are novel and hard to predict but from which nothing meaningful can be learned (e.g., the pattern of raindrops falling on the ground). Instead learners, including human infants, set their own goals for learning, selectively attend to information that is learnable, and decide what and whom to learn from (e.g. Begus, Gliga, & Southgate, 2016; Gerken, Balcomb, & Minton, 2011; Kidd, Piantadosi, & Aslin, 2012). Although much remains to be understood about how these goals are established and constrain learning (see discussion to follow), there is broad consensus that learners are most motivated to explore when there is neither too much information to be learned nor too little (Begus et al., 2016; Berlyne, 1960; Csikszentmihalyi & Csikzentmihaly, 1990; Dember & Earl, 1957; Gerken et al., 2011; Gottlieb et al., 2013; Kidd et al., 2012; Kinney & Kagan, 1976; Loewenstein, 1994; Oudeyer, Kaplan, & Hafner, 2007; Schmidhuber, 2013) – a "goldilocks effect" that has been attributed variously to the learners' representation of the surprisal value of stimuli (Kidd et al., 2012), the rate of change in their own learning (Gottlieb et al., 2013), the size of the gap in their current knowledge (Loewenstein, 1994), or the amount of structure in the stimuli (Gerken et al., 2011).

finally, in recent years, there has been substantial progress in thinking about the computational and neural substrates that might subserve effective information seeking, especially in contexts where rewards are sparse (Burda et al., 2019; Chitnis, Silver, Tenenbaum, Kaelbling, & Lozano-Pérez, 2021; Gottlieb et al., 2013; Oudeyer, Gottlieb, & Lopes, 2016; Oudeyer et al., 2007; Pathak, Agrawal, Efros, & Darrell, 2017; Schmidhuber, 2013). The various approaches differ both between and within fields but they share a commitment to curiosity-driven exploration as a means of gaining information, reducing uncertainty, and improving prediction and control. In this sense, all these accounts attest to the idea of "play for prediction and plans".

2.4 Distinctively human play

These five accounts – play for pleasure, play for performance, play for peace-making, play for practice, and play for prediction and plans – each cover a broad range of behavior in themselves, and as discussed, are not mutually exclusive. A six-month-old baby might grasp a rattle for the pleasure of holding something tightly in his hand and shake it vigorously, conveying health and fitness. The sight and sound might amuse others and strengthen social bonds. In exercising his fine-motor development, the play might make him a better tool-user as an adult. And by coming to anticipate the sound made by his shaking, he might develop better predictive models of the world and be able to organize his own behavior into increasingly complex sequences.

In short, even if no one of these accounts does justice to the richness of play, collectively, they might seem fairly comprehensive. But the limitations of these accounts, and the degree to which play remains elusive, may become more evident when, thirty-months later, that same child, armed with a kitchen strainer, takes that rattle, buries it in a hole, covers it with leaves, and when asked what he is up to, explains that he is building a trap for a velociraptor because when the velociraptor steps on the leaves, the rattle will make a noise, and then he can trap it with the strainer.

The example is frivolous but the point is not. Although the opportunities, content, and resources available for children's play vary cross-culturally, the richness and variability of play is a human universal (Edwards, 2000; Gosso, 2010; Gosso, e Morais, & Otta, 2007; Lancy, 2002, 2007; Nwokah & Ikekeonwu, 1998; Schwartzman, 1986; Singer, Singer, D'Agnostino, & DeLong, 2009). The kind of elaborated behavior, seamlessly integrating elements of pretend and exploratory play [4], depicted in the

[4] In the discussion to follow, we will draw no distinction between exploratory and pretend play largely because it is not obvious that children do. A child pretending to trap a velociraptor may explore whether the rattle can be used to dig a hole in the leaves; a child playing with stacking cups may be making a dinner table for his toy hippopotamus. For our purposes, what different forms of distinctively human play have in common will be more important than what distinguishes them. Also, from here forward, unless otherwise specified, "play" will serve as shorthand for "distinctively

velociraptor example is not the exception but the rule in human play after toddlerhood.

And critically, there is no sense in which play like this is merely a generalization of the kind of play characterized by the accounts above. This kind of play is often solitary, so it is unlikely to be useful for performance or peacemaking; nor does it seem primarily driven by a pleasurable sensorimotor component. One could suppose that the child is practicing adult activities like, say, trapping mice but it stretches credulity to suppose that this kind of play will actually help him build a better mousetrap. And this kind of play defeats even the most cognitively sophisticated account we reviewed: play for prediction and planning. The child already knows the rattle makes noise. He also already knows that the leaf pile will conceal the rattle: that is why he's hiding it in there. To play the way he is, he already had to access abstract concepts of concealment, detection, and containment and coordinate them into complex plans. And needless to say, the activity is unlikely to teach him anything he didn't already know about velociraptors.

So what is the child doing? And what, if anything, does it have to do with distinctively human curiosity and cognition? The remainder of this chapter is an attempt to answer these questions. To foreshadow, we will suggest that distinctively human play involves manipulating our own utilities such that we invent problems for ourselves. We will offer some preliminary evidence for this, showing that even given identical goals, children's exploratory behavior differs from their exploratory play; children violate normal utilities only when they are playing. We will then briefly digress from play to explain why problems, in general, might be valuable for hypothesis generation. We will then return to the topic to propose that the sustained engagement children show in play, independent of any obvious reward (including information gain), is a hallmark of distinctively human curiosity: a curiosity that depends not on progress in learning but in thinking. finally, we will suggest that the idiosyncratic, arbitrary nature of the problems set in play, and the often flimsy, inadequate solutions generated, may be offset by the fact that the ideas generated in play can be decoupled from the problems that inspired them and be valuable in their own right. Throughout this section, the account we offer is speculative, relying more on conjectures than data. But although we approach the topic somewhat playfully ourselves, we do so with the serious intent of trying to grapple with what is distinctive about human play, curiosity,

human play of the kind that emerges after toddlerhood". But nothing that follows should be taken as disputing the prevalence of other kinds of play or invalidating any of the ways of accounting for those kinds of play reviewed above.

and cognition.

2.4.1 Play for problems and proposals

As described above, scientists can use play as a dependent measure of children's sensitivity to many factors connected to information gain: violations of intuitive theories (Bonawitz et al., 2012; C. Legare, 2014; Schulz, Standing, & Bonawitz, 2008), the information structure of tasks (Ruggeri, Swaboda, Sim, & Gopnik, 2019), the ambiguity of hypotheses (Cook et al., 2011; van Schijndel et al., 2015), and the discriminability of data (Siegel et al., 2021). Arguably however, these tasks tell us a great deal about children's *learning* but relatively little about children's *play*. That is, although we, as adult scientists, can use play to assess children's sensitivity to uncertainty and expected information gain, that is not necessarily the best characterization of what children use it for.

Some prima facie evidence that this is the case is the notorious gap between play as we study it in the lab and play as it exists in the wild. It takes considerable effort and bespoke experimental designs to contain the variability and arbitrariness of children's play sufficiently for it to be used as a measure of information gain. In work on play as a form of rational learning, that variability is treated as noise. But children's propensity to adopt idiosyncratic goals may be what distinctively human play is all about. More broadly, we propose that many kinds of distinctively human play – from catching velociraptors to building rocket ships to playing soccer or chess – involve creating problems for ourselves. We suggest that the most salient characteristic linking all these forms of play is the extent to which we intervene on our own utility functions and willingly incur unnecessary costs to achieve arbitrary rewards.

2.4.2 Decoupling utilities from utilitarian ends

The degree to which play involves manipulated utility functions can perhaps be best seen in directly comparing children's exploration with their exploratory play. Although much of the literature on play and learning (including most of the senior author's own work) has treated these as equivalent behaviors (i.e., play as rational exploration) we believe this is misleading. Specifically, in recent work (Chapter 4, 2020a), we gave four and five-year-old children closely matched retrieval and exploration tasks which differed only in whether children were asked to achieve a goal or asked to play and achieve the goal. Thus for instance, children were brought to the door of a room with

Figure 2-4: Play for problems and proposals: in the studies presented in Chapter 4 (Chu & Schulz, 2020a), preschoolers were told to play and retrieve objects. In play children (a) reached for out-of-reach pencils and ignored the easily-accessible pencils in a cup on the table, and (b) walked in a spiral to get to a box of stickers in the middle of the room rather than running straight for it. When asked to retrieve the same objects for instrumental reasons, children did the opposite: going straight to the pencils in the cup and the box of stickers.

a spiral design on the floor and a box at the center of the spiral. In one condition, the children were told "There are stickers in that box. Can you go in here and try to get one?" In the other, they were told "There are stickers in that box. Can you play in here and try to get one?" In the former case, children ignored a spiral design on the floor and walked in a straight line to the box in the middle of the spiral; in the latter, they not only walked around the spiral before getting the stickers but sometimes did so twice or walked around backwards (see **Figure 2-4a**). In a different experiment, children were introduced to a room containing a table with a cup of pencils on it and a stenciled tree on the wall with pencils velcroed to the branches just out of the children's reach. When told, "I need a pencil to fill out this form. Can you go over there and try to get a pencil?" children went directly to the cup. But when told "I need to fill out this form. While I'm doing that can you play over there and try to get a pencil?" children went to the stenciled tree and not only jumped up to get a pencil, but having retrieved one pencil, then jumped up and down again repeatedly to try to get more pencils near the top of the tree (See **Figure 2-4b**). Children showed the same pattern in exploration tasks. Given a choice between one drawer on the left and twelve drawers on the right, children reliably preferred the smaller search space when they were told to find a ball to use in another game and reliably preferred the larger when told it was a hide and seek game ("I'm going to hide the ball and you get to find it. Do you want to play over here or over there?").

In the sense that human play involves manipulated utility functions, all play is pretend play; when the costs or rewards are real, we are no longer playing. And yet, it would be a mistake to think that in play, children behave either randomly or irrationally. Even when children opted for the unnecessarily costly goal, they behaved efficiently with respect to that goal (adhering close to the spiral path; jumping directly towards the pencils; searching the boxes in sequential order). Thus we suggest that in play, children's behavior is not only boundedly rational (e.g. limited by children's information processing constraints; Simon 1955) but *conditionally rational*: rational with respect to a manipulated utility function.

Of course, humans are not the only animals to engage in self-handicapping behavior during play. Primates will try to balance themselves on unstable branches and wolves and dogs will bow low rather than towering to attack in play fighting. Researchers have proposed that this kind of play prepares animals for unusual, unexpected events and may provide them with experience improvising solutions (see Spinka et al., 2001). We suggest that a violation of normal utility functions characterizes human play as well but in much more far-reaching ways. We not only incur unnecessary costs, but also flexibly fix our own arbitrary rewards by setting a vast range of novel goals.

2.4.3 How problems structure their own solutions

But surely there are enough problems in the world. Why should we make new ones for ourselves? We believe novel problems and goals[5] may be critical to human cognition because problems constrain search, and narrowing the search space sufficiently to generate new hypotheses is arguably, far more than learning per se, the hard problem of cognition. To quote from a recent workshop on program induction "Coming up with the right hypotheses and theories in the first place is often much harder than ruling among them ... How do people, and how can machines, expand their hypothesis spaces to generate wholly new ideas, plans, and solutions?" (Bramley, Schulz, Xu, & Tenenbaum, 2018).

The idea that goals could improve learning and planning is widely recognized in AI and robotics. Many approaches to engineering intrinsically motivated autonomous

[5]We will use "problems" and "goals" interchangeably here on the understanding that if you have a goal, then the problem is how to achieve it. Problems and goals differ chiefly in that when your problem is a query your solution can take the form of a hypothesis or proposal; when your problem is how to achieve a goal, the solution must take the form of a plan. For our purposes, the key point is that in both cases, the information in the problem itself constrains the search for solutions.

agents involve having agents establish their own curricula for goal-directed learning (e.g. Agrawal, Nair, Abbeel, Malik, & Levine, 2016; Chitnis et al., 2021; Kaelbling, 1993; Lynch et al., 2020; Sukhbaatar, Kostrikov, Szlam, & Fergus, 2017). However, the current proposal differs from these accounts in our commitment to the idea that the goals generated in play might not translate into reduced uncertainty or prediction error, or even the achievement of the goals themselves. (Indeed, there's not any obvious sense of what it might *mean* for a child to successfully trap an imaginary velociraptor.) Rather, the value of the problems posed in play might be simply in generating the new thoughts and plans themselves. Liberated from any real-world goal – even the goal of fulfilling its own goals – human play may be not so much a means of gaining information as a means of increasing innovation.

But to understand how problems contain the kind of information that could support the generation of new ideas, we must leave play behind for a moment and turn to problems themselves. To start, consider the information available in question words (see **Figure 2-5**.)

Before you know what someone is asking, let alone be able to answer their question, you already know a lot about what the answer has to look like. Answers to "who" questions are likely to refer to a social network; answers to "where" questions to a map; "when" questions to a timeline; "what" questions to a category structure; "which" questions to the intersection of a venn diagram, "how" questions to a circuit of some kind, "why" questions to a causal network. And each additional word in a query adds further constraints, further narrowing the search space for the solution. Even just a single additional function word can do a lot of work: "Why does . . . " will likely be answered by a rule or empirical generalization; "Why did . . . " will have to account for an unexpected event; "Why can't . . . " will have to explain why something seemingly possible or permissible is not. Each word of the query imposes different constraints on the possible responses.

Knowing the form of a question doesn't mean you can answer it, but at least it will get you in the right ballpark. Rejecting wrong answers to a question may be far more tractable than evaluating answers that are not even wrong. (e.g., "1774" is the wrong answer to when the United States declared independence from Great Britain but it is "not even wrong" in response to the question of *why* the United States declared independence from Great Britain).

While representing the abstract space of answers to a query might seem to require

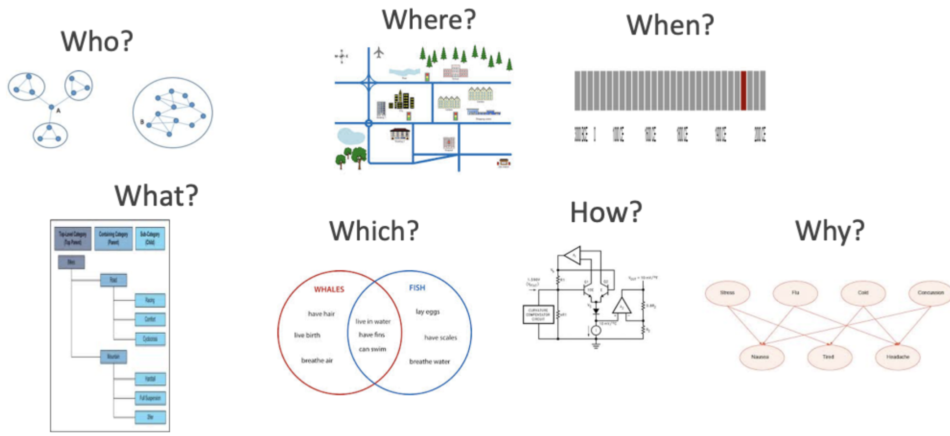

Figure 2-5: Queries contain information constraining their own answers, independent of content-domain.

sophisticated reasoning, as early as two and three, young children are sensitive to the form of questions and what might count as answers (Bloom, Merkin, & Wootten, 1982; Callanan & Oakes, 1992; Frazier, Gelman, & Wellman, 2009, 2016; Nelson, Divjak, Gudmundsdottir, Martignon, & Meder, 2014). And despite preschoolers' robust preference for reliable, confident informants (e.g. Harris, Koenig, Corriveau, & Jaswal, 2018; Jaswal & Malone, 2007; Koenig, Clement, & Harris, 2004), children will accept tentatively advanced conjectures that are possible answers to a question over confidently asserted, known facts that are not (Chapter 3, (Chu & Schulz, 2021)). Thus, simply *posing* a question might allow learners to start generating plausible solutions (see Schulz, 2012a for discussion).

And critically, rich, structured constraints are not a property only of queries expressed in language: all kinds of problems are rich in information. To the degree that we can represent abstract properties of our problems and goals (e.g., "I need something that is smaller than this but the same shape" "I want this to go up and down again and again" "I need something that gets bigger fast") these representations (e.g., proportionality, cyclic variation, non-linear growth) could constrain the search for solutions and plans, independent of content-domain. Previous work suggests that four and five-year-olds can use abstract properties like proportionality and cylicity to constrain hypotheses about probable causes of observed effects (Magid, Sheskin, & Schulz, 2015; Tsividis, Tenenbaum, & Schulz, 2015). Thus abstract representations of the form of problems, together with their specific content, might provide sufficient constraints for children to generate new ideas and plans.

2.4.4 Problems and distinctively human curiosity

We will dodge (for lack of an answer) the question of how learners recognize when problems have sufficient structure to be tractable, and simply suggest that humans are sensitive to the extent to which our problems constrain the search for solutions[6]. We propose that our recognition that a problem is tractable – in the sense of containing enough information to guide the search for a solution – inspires the kind of curiosity which can sustain long-term engagement in the face of negligible information gain.

We can contrast this with a large body of work on epistemic curiosity suggesting that actions with high expected information gain are reinforced (or not) with respect to the degree to which they reduce online uncertainty and prediction errors. This kind of curiosity motivates human exploration and learning in any number of paradigms, including visual search tasks (Gottlieb et al., 2013; Kidd et al., 2012), bandit tasks (Daw, O'doherty, Dayan, Seymour, & Dolan, 2006), seeking information about risky choices (Blanchard et al., 2015; Bromberg-Martin & Hikosaka, 2009, 2011), opening doors to reveal hidden objects (Jirout & Klahr, 2012), learning the answer to trivia questions (Loewenstein, 1994), and autonomous artificial agents' exploration of novel spaces and novel objects (e.g. Burda et al., 2019; Florensa, Held, Geng, & Abbeel, 2017; Forestier, Mollard, & Oudeyer, 2017; Friston et al., 2015; Little & Sommer, 2013; Martius, Der, & Ay, 2013; Oudeyer et al., 2016; Pathak et al., 2017; Schmidhuber, 2010). Indeed, the virtue of such accounts is that they account for exploration broadly and extend to the kinds of epistemic curiosity that might apply across many intelligent agents.

But if epistemic curiosity tracks our progress in *learning* (e.g., Oudeyer et al., 2016, 2007; Schmidhuber, 1991) – the degree to which our predictions improve and our uncertainty decreases – it is hard to explain the kind of sustained engagement that characterizes much of our experience as humans. Humans can be fascinated both by questions we will never answer: ("Who would you be if you had a brain transplant?" "What would you do if you had a billion dollars?") and by questions that may take years or even lifetimes to answer ("Can we find particles of dark matter?"). We suggest that this kind of epistemic curiosity is consistent with an ability to track, not only the rate at which we are *learning* but the rate at which we are *thinking*. To the degree that we can continue to generate speculations, hypotheses and partial plans we may feel like we are making progress on a problem even if there is no evidence to assess

[6]And as you might have observed, we will also dodge (again, because we don't have an answer) the question of how learners generate new goals and problems.

that progress. The degree to which a problem or goal supports the generation of plans and hypotheses may itself be motivating – whether or not the plans actually bring us closer to attaining the goal and whether or not those hypotheses reduce prediction error. That is, we propose that in humans, intrinsic reward is tied to the ability to act and and think, not merely the consequences of our actions and thoughts [7].

Precisely because this kind of engagement is no guarantee of increased learning, it is, quintessentially "playful". And indeed, we believe children's often rapt absorption in solitary play suggests the intrinsic reward associated merely with thinking. Of course, there's a very large gap between inventing plans to catch non-existent velociraptors and inventing plans to catch possibly existent weakly interacting massive particles. But humans are the only creatures that do both, and to be kinds of creatures we are – learners whose learning goes far beyond prior knowledge and the data – we may have to value problems that engage us merely to the extent that we can generate possible solutions. Actual solutions – ones that reduce uncertainty and increase prediction and control – can come only later, if at all.

2.4.5 Unknown unknowns and exploring new ways to explore

Thus there is a sense in which we must take the uselessness of play seriously. The problems invented in play are arbitrary and unimportant (try to trap a velociraptor; feed dinner to the hippo; get a ball through a net; use black or white stones to surround space on a grid, etc.). What seems true at face value – that burying rattles in leaves or building dinner tables for hippos serves no purpose – is, we suggest, perfectly true if by that we mean that these activities neither prepare children for adult life nor reduce children's uncertainty about anything they were uncertain about. However, although this kind of distinctively human play may be useless for many ends, we have argued that it may be useful for thinking; even frivolous problems contain enough structure and information to allow us to start generating new thoughts and plans. Still, one might wonder, what is the use of thinking frivolous things? Especially when, by many criteria, the ad hoc solutions we generate in response to arbitrary problems we invent

[7] These ideas are loosely connected to ideas about "empowerment" in the reinforcement learning literature – the idea that organisms are intrinsically motivated to maximize the degrees of freedom they have for acting on their environment (Klyubin, Polani, & Nehaniv, 2005; Salge, Glackin, & Polani, 2013). However, empowerment indexes the extent to which an organism can influence its environment and register its influence. The current account by contrast, is not predicated on the idea of control but on the capacity for thought and action itself (independent of any downstream effects of those thoughts and actions). That is, we suggest that the intrinsic value of a goal is its ability to generate a plan; the intrinsic value of a problem is its ability to generate possible solutions.

are themselves bad ones (cf: the plan for trapping the velociraptor).

One reason that the triviality of the problems and the badness of our proposed solutions may not matter is that the ideas we generate can be decoupled from the problems that inspired them. Indeed one of the striking and characteristic features of children at play is that they often spend much of the day playing and then abandon their plans both without ever achieving their goals without any apparent regret. When a seven-year-old decides to build a spaceship and fly to Mars, she may have very decided opinions about what to do and how to do it, spend hours tinkering with tinfoil, tape, and your hairdryer, but then abandon whole setup in the backyard after half a day's work without a flicker of dismay the instant the icecream truck rolls by.

As discussed, we suggest that what matters about the child's play is neither the unachievable goal nor the half-baked solution but the fact that the goal contained just enough structure to generate new ideas. The little that the seven-year-old knew about rocket ships (e.g., that they are large, shiny and propelled by something) was sufficient to support thinking and planning. But although a hairdryer is a very bad way to move a rocketship, it may not be a bad way to make something move in general. The idea can be decoupled from the goal that motivated it and may just possibly turn out to be useful in unrelated contexts (e.g., for extracting a retainer from under a bookshelf).

Critically, it is not the case (on our account) that the child *learned* that hairdryers could propel things through play. If it hadn't occurred to the child that the hair dryer could propel things, she would never have swiped it from her mother's drawer. Rather, the point is that the child would have been unlikely to *think* about the propulsive capabilities of hair dryers but for the fact that she wanted to build a rocketship and needed something that would serve that purpose. Posing the problem she did supported thinking of the solution she did. Once thought of, the solution can take on a life of its own. Thus the idiosyncratic, arbitrary nature of play may be valuable because each new problem and goal imposes unique constraints, which lead to unique plans and solutions – any of which may be repurposed. One reason our motivational system may be as rich as it is, is because our ability, as a species, to want anything at all lets us explore a vast space of possible plans and ideas.

More broadly, as we have reviewed, all kinds of animals engage in exploration and learning. What might be distinctive about human play is that people do not just exploit their existing knowledge about how to explore (i.e., by acting efficiently to maximize expected information gain) but instead, explore new ways to explore. A sure way to do that is to intervene on normal utility functions, assigning arbitrary rewards and accepting unnecessary costs. But again, *why* explore new ways to explore? Why not just explore in ways most likely to increase learning? Arguably because epistemic goals are not the only — or necessarily even the best — way to learn new things. The world is full of unknown unknowns; as great as our uncertainty about the world is, there are even more things we don't even know we don't know. If we only explored to try to maximize expected information gain, we would miss the chance to gain unexpected information. Creating new problems with no obvious utility in themselves – playing – may be the best way to discover (genuinely) new things.

2.5 Conclusions and future questions

We have suggested that in addition to play being valuable for pleasure, performance, peace-making, practice, and prediction, human play may be valuable in supporting the creation of new problems and goals; these in turn, may support new thoughts, plans, and discoveries. However, at the moment, this account is just a speculation; many questions remain unanswered. As we've noted, goals and problems might support search by constraining the hypothesis space, but we've deferred the question of how we generate new goals and problems themselves. What is it about human minds that make our utilities so flexible such that we can assign value to almost anything and willingly incur unnecessary costs? How do we distinguish ill-posed problems which insufficiently constrain search from those that are rich in structure and therefore potentially tractable? Can we formalize the information in problems and goals well enough to specify how they support the generation of new thoughts and plans? And how do we represent our own progress in thinking such that it can be a source of intrinsic reward? These and many other questions ensure that we are likely to remain curious about play for years to come.

We will end by noting that scientists have puzzled over the relationship between play and development for well over a century (indeed, the philosopher and psychologist Karl Groos' 1898 formulation of "play for practice" inspired the alliterative trope here), and some of us, personally, have puzzled over the relationship between play and learning for most of our adult life. Our enduring fascination with this topic is itself a source of some mystery. Science is supposed to answer questions. Surely there should be something dismaying about finding that the questions persist? But as we have suggested here, we may be most curious, not to the degree that we anticipate being able to answer our questions but to the degree that we realize we may never stop thinking about them. This as much as anything may be the signature of a human mind at play.

Speculations are just ideas - no data or experiment design in this proposal yet. How do we decide if a novel proposal is a good idea when we don't have data yet? In the next chapter, we explore how adults and young children assess speculative conjectures in the absence of evidence.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Children selectively endorse speculative conjectures

> Bold ideas, unjustified anticipations, and speculative thought, are our only means for interpreting nature
>
> Popper
> The Logic of Scientific Discovery

Abstract

Young children are epistemically vigilant, attending to the reliability, expertise and confidence of their informants and the prior probability and verifiability of their claims. But the pre-eminent requirement of any hypothesis is that it provides a potential solution to the question at hand. Given questions with no known answer, the ability to selectively adopt new, unverified, speculative proposals may be critical to learning. This study explores the conditions under which people might reasonably reject known facts in favor of unverified conjectures. Across four experiments, when conjectures answer questions that available facts do not, both adults (n=48) and children (4.0-7.9 years, n=241) prefer the conjectures, even when the conjectures are preceded by uncertainty markers or explicitly violate prior expectations.

3.1 Introduction

The history of science is full of remarkable discoveries but it is also renowned for speculative conjectures that were ultimately abandoned. Scientists posited the existence of a luminiferous aether to explain how light traveled through a vacuum, phlogiston to explain the behavior of fire, and elan vital to explain life itself. All of these ideas were ultimately discredited but today, other suppositions that have not been directly confirmed (e.g., the existence of dark matter and dark energy) play powerful explanatory roles. Faced with otherwise unanswered questions, we must sometimes entertain claims whose primary value lies, not in how certain we are of their truth, but in how certain we are that – if they were true – they would provide a solution to our problems.

But our willingness to entertain potentially explanatory but unproven speculations extends far beyond the realm of scientific inquiry. The comedian Stephen Colbert mocked politicians' readiness to accept plausible but unsubstantiated arguments as a predilection for "truthiness" over "truth" (cf: at the time of this writing, the conjecture that disinfectants might kill COVID-19 inside our bodies as well as out). However, we suggest that our ability to accept proposals based merely on the possibility that they could answer a question is not (just) a bug but also a feature of human cognition. Conjectures go beyond available knowledge and data but they are not entirely unconstrained; even the wildest speculation must, in principle if not in practice, provide an answer to the question under discussion. Insofar as one of the most challenging problems of human cognition is not distinguishing among competing hypotheses but generating them in the first place, even proposals advanced without evidence, by uncertain speakers, and in tension with prior expectations may be worth considering if they offer possible solutions to otherwise unresolved problems. Further investigation can then establish whether the hypotheses should be pursued or rejected.

We suggest that a willingness to entertain claims merely on the basis of their explanatory power is a pervasive aspect of human cognition, beginning in very early childhood. However, while many studies have looked at how children evaluate both the quality of the explanations they receive and the reliability of their informants, such work has focused almost uniformly on whether children correctly reject improbable, unreliable, or unsubstantiated claims in favor of information that is trustworthy, verified, and consistent with the integration of evidence and prior knowledge. To our knowledge, no studies have looked at situations in which either adults or children might (appropriately) reject known information in favor of the unknown. Here we look at the conditions under which people might reasonably endorse conjectural claims.

Contrary to early assumptions about the credulity of children (Clark, 1990; Piaget, 1929; Prentice & Gordon, 1987; Prentice, Manosevitz, & Hubbs, 1978), and despite the importance of pretense and fantasy in children's lives (Harris, 2000; Lillard, 2001; Lillard, Pinkham, & Smith, 2011; Sharon & Woolley, 2004; Taylor, Cartwright, & Carlson, 1993; Walker, Gopnik, & Ganea, 2015; Weisberg & Gopnik, 2013; Woolley, 1995), even three and four-year-olds do not adopt fanciful, speculative claims willy nilly (see e.g., Harris et al., 2018; Ronfard, Zambrana, Hermansen, & Kelemen, 2018; Shafto, Eaves, Navarro, & Perfors, 2012; D. Sobel & Kushnir, 2013; Sperber et al., 2010; Woolley & Ghossainy, 2013). Children evaluate the reliability of their informants in increasingly sophisticated ways from preschool through middle childhood (Clément, Koenig, & Harris, 2004; Gweon et al., 2014; Koenig et al., 2004; Koenig & Harris, 2005; Koenig & Woodward, 2010; Pasquini, Corriveau, Koenig, & Harris, 2007). They are sensitive to the strength of the evidence they observe (Bridgers, Buchsbaum, Seiver, Griffiths, & Gopnik, 2016; Butler, Schmidt, Tavassolie, & Gibbs, 2018), the prior probability of testimony (Chan & Tardif, 2013; Clément et al., 2004; Jaswal, 2004; Koenig et al., 2004; Koenig & Echols, 2003; L. Ma & Ganea, 2010), the informant's past accuracy, knowledge and expertise (Danovitch & Keil, 2004; Koenig et al., 2004; Koenig & Harris, 2005; Koenig & Jaswal, 2011; Kushnir, Vredenburgh, & Schneider, 2013; Landrum, Mills, & Johnston, 2013; Nguyen, 2012; D. Sobel & Corriveau, 2010; D. Sobel & Macris, 2013; VanderBorght & Jaswal, 2009), and the situational and epistemic constraints the informant is under (Butler et al., 2018; Einav & Robinson, 2011; Flavell, 1988; Nurmsoo & Robinson, 2009; Senju, Southgate, Snape, Leonard, & Csibra, 2011).

This does not mean that all of children's evaluation of informants is epistemically justified. Preschoolers have a general bias in favor of agents who are friendly (Brosseau-Liard & Birch, 2010; Landrum et al., 2013), familiar (Reyes-Jaquez & Echols, 2013), attractive (Bascandziev & Harris, 2014; Fusaro, Corriveau, & Harris, 2011), members of their in-group (Elashi & Mills, 2014; Kinzler, Corriveau, & Harris, 2011; MacDonald, Schug, Chase, & Barth, 2013; Plötner, Over, Carpenter, & Tomasello, 2015; Wood, Kendal, & Flynn, 2013), or part of a majority (Chen, Corriveau, & Harris, 2012; Corriveau & Harris, 2010; DiYanni, Corriveau, Kurkul, Nasrini, & Nini, 2015; Morgan, Laland, & Harris, 2014). Children are also influenced by the confidence with which informants assert their claims. Preschoolers are more likely to endorse novel explanations and labels advanced with confidence than those provided by a speaker who is hesitant or expresses uncertainty (Jaswal & Malone, 2007; Kominsky, Langthorne, & Keil, 2016; Sabbagh & Baldwin, 2001; Tenney, Small, Kondrad, Jaswal,

& Spellman, 2011). However, even these prima facie non-epistemic biases may be reasonable routes to learning insofar as friendly, familiar, in-group members who are backed by a majority consensus may typically also be the most likely sources of reliable information. Collectively, this literature suggests that children might be very likely to reject information that is unverified or unverifiable, especially if the speaker conveys uncertainty or the information is itself unlikely or unexpected. However, children are also interested in getting answers to their questions. Children are notorious for asking questions themselves (as many as 76 an hour; Chouinard 2007), and although some are requests for permission, or for information redundant with facts the child already knows (C. Legare, Mills, Souza, Plummer, & Yasskin, 2013; Ruggeri, Lombrozo, Griffiths, & Xu, 2016), many are requests for novel information and explanations (Callanan & Oakes, 1992; Chouinard, 2007; C. M. Mills, Legare, Bills, & Mejias, 2010; C. M. Mills, Legare, Grant, & Landrum, 2011).

Preschoolers also understand many structural aspects of explanation and can evaluate respondents' answers on those bases alone. If a respondent simply restates a child's question, asserts norms, re-describes events, or reacts personally instead of responding to the query, the child is likely to repeat the question (Chouinard, 2007; Frazier et al., 2009; Kurkul & Corriveau, 2017; Tizard & Hughes, 1984). Preschoolers favor claims supported by strong arguments over circular ones, and circular arguments over unsupported opinions (Corriveau & Kurkul, 2014; Mercier, Bernard, & Clement, 2014; C. M. Mills, Danovitch, Rowles, & Campbell, 2017) and evaluate explanations based on how many observations an explanation accounts for, how simple and internally coherent it is, and how probable it is given observed data and their prior knowledge (Bonawitz & Lombrozo, 2012; Johnston, Johnson, Koven, & Keil, 2016; Lombrozo, 2011; Walker, Bonawitz, & Lombrozo, 2017).

Children can also use data-independent criteria to evaluate hypotheses that lack direct evidential support. In addition to favoring explanations that are simple, broad and coherent (Bonawitz et al., 2012), children can use properties of the explanandum when choosing among equally probable hypotheses. When asked to match observed events to their probable causes, preschoolers expected discrete and continuous affordances to control discrete and continuous phenomena respectively, without observing any covariation data (Magid et al., 2015; Tsividis et al., 2015). This suggests that children might be sensitive to abstract features of causes and effects and use these features to constrain their generation and evaluation of candidate causes. Independent of the content of the domain, and in the absence of any distinguishing evidence, children might be able to use properties of the question under discussion to decide what makes for a good answer.

In the current studies, we look at whether children flexibly evaluate facts and conjectures given questions that can or cannot be answered by available information. In Experiment 1 we introduce children to short stories involving novel characters and events. We ask children to choose between factual and conjectural explanations, for questions that can or cannot be answered by information in the story. We used nonsense characters and stories in order to control for effects of prior knowledge, and we matched the answers on the degree to which they repeated the words from the question in the answer. Thus, children must consider the degree to which each response provides a potential answer to the question at hand. In Experiment 1a, we tested a wide age-range (four to eight-year-olds) given that it was not clear to what extent children at different ages might privilege abstract features of explanations over established facts.

If children always prefer the most certain and reliable information, they should always choose the facts; if they always prefer more speculative, inventive answers, they should always choose the conjectures. However, we predict that the six to eight-year-olds, and possibly even the preschoolers, would prefer the known facts for questions that can be answered by information in the story and prefer conjectures for questions that cannot.

3.2 Experiment 1

We ran both an initial exploratory study (Experiment 1a) and a replication with just the four and five-year-old's (Experiment 1b). Hypotheses were prespecified ahead of data collection, but not formally preregistered. For the initial experiment, we estimated a moderately large effect size in choosing facts for questions with available answers and conjectures for questions with unknown answers. We aimed to recruit 64 participants in Experiment 1a which would yield 80% power to detect an odds ratio of 5.23 (pilot testing had suggested an odds ratio of 6.93). In Experiment 1b, we tested only a younger age group, and recruited 32 participants to match the number of younger children in Experiment 1a.

3.2.1 Methods

Participants

All children in this and the following experiments were recruited from an urban children's museum between January 2018 and November 2019 in the United States. Parents provided informed consent, and children received stickers for their participation. Although most of the children were white and middle class, a range of ethnicities and socioeconomic backgrounds are represented in museum attendees overall (47% European American, 24% African American, 9% Asian, 17% Latino, 4% two or more races; 29% of museum attendees visit on days when there is free or discounted admission).

In Experiment 1a, we tested 66 children, ages four to eight (M = 6.04 years, range: 4 – 7.93). Seven additional children participated but were excluded for either responding inaccurately on a practice question (N = 5), not speaking English as their primary language (N = 1) or for incomplete participation (N = 1).

In Experiment 1b, we tested 32 four- and five-year-olds (M = 5.03 years, range: 4.15-5.92). Thirteen additional children did not pass the inclusion criteria (9 failed practice; 2 did not speak English as their primary language; 1 withdrew; 1 did not respond to test questions). The exclusion rates for four and five-year-olds are relatively high but can be explained almost entirely by children choosing Elmo as the correct puppet on both practice trials.

Materials and procedure

Each trial began with an illustrated story presented via three animated slides on a laptop computer. See **Figure 3-1** for an example story. Two puppets (Elmo and Cookie Monster) were also used; the puppets sat on either side of the computer and "watched" the stories with the child. The puppets' answers were delivered by pre-recorded audio clips to avoid inadvertently biasing the children with prosodic cues.

Participants completed two training trials and four test trials. The training trials were designed to ensure that participants were paying attention and understood the task. Participants who failed the training trials were excluded from analysis and replaced. These stories depicted human characters performing common activities (i.e., riding a bike; eating ice cream) embedded in a simple narrative. The training questions could always be answered using information from the story. Each puppet provided a correct answer on one trial and an incorrect trial on the other (order counterbalanced).

Children were tested individually in a quiet room. The experimenter began by introducing participants to the computer display and the puppets (Elmo on the left and Cookie Monster on the right). The experimenter explained the task, saying: "*Every time I tell you a story, I need you to remember what happened because I'm going to ask a question at the end. Elmo and Cookie Monster will tell us their answers and your job is to choose who had the better answer.*" On every trial, the experimenter first narrated the story and presented her question ("My question is, ...?"). She then directed the question at one puppet (e.g. "*Elmo, can you tell us, ... ?*"), played its pre-recorded answer, and repeated the answer (e.g. "*Elmo said because ...*"). The experimenter then repeated the question-answer sequence with the other puppet before repeating the question and inviting the child to make a choice ("*My question was ... Who do you think had the better answer for [question]?*"). Positive feedback was given on the training trials ("*That's right, Elmo had the better answer this time.*") and neutral, encouraging feedback was given on the test trials ("*Alright, let's see what's next*"). Only children who correctly answered both practice questions continued to complete the four test trials.

The test trials involved imaginary creatures engaging in different activities (making a hat, sneezing from allergies, dropping a toy down a deep hole, juggling). Two question-answer pairs were used on each story: one question could not be answered with the conjecture offered but could be answered with a Fact mentioned in the story (In Story Question); the other question could not be answered with any facts in the story and could only be answered with a Conjecture (Out of Story question). Regardless of question type, Elmo always provided the Fact answer and Cookie Monster always provided the Conjecture answer. Elmo always provided his answer first. To cover a range of explanatory question types, test trials included both "why" and "how" questions. In Experiment 1a, two "how" questions came first; in Experiment 1b, two "why" questions came first. In both experiments we counterbalanced two between-participant factors: (1) item order (whether the first test trial was an In-Story or Out-of-Story question) and (2) story-question match, resulting in four story sequences. Thus, while all participants heard all stories and answers, half the participants heard any given story presented with a question that could only be answered with a fact and half heard the story presented with a question that could be answered only with a conjecture. See Table 3.1.

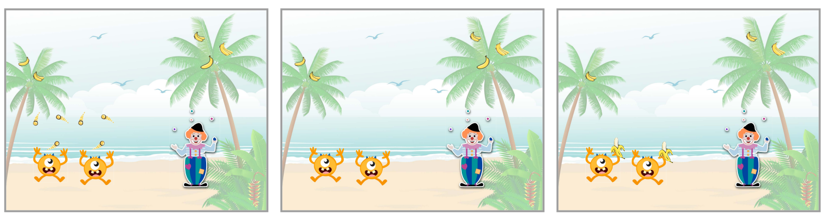

[Slide 1] Here are some juggling Gazzers. A clown named Bozo taught them to juggle.

[Slide 2] Juggling Gazzers love to eat bananas. But the bananas all grow at the top of very tall trees and the Gazzers can't climb trees.

[Slide 3] But here the Gazzers are! Eating bananas.

In Story Question: How did the Gazzers learn to juggle?

Out of Story Question: How did the Gazzers get the bananas?

Fact: Because Bozo the clown taught the banana eating Gazzers how to juggle

Conjecture: Because the Gazzers threw their balls up into the trees and knocked down the bananas.

Figure 3-1: Example of a test trial used in Experiments 1 and 2.

Table 3.1: Questions and candidate answers used in Experiments 1 and 2.

Story	**In-Story Question**	**Out-of-Story Question**	**Fact Answer**	**Conjecture Answer**
Training 1	How did Tommy get to the castle?	-	He rode a bike to the castle	He walked to the castle
Training 2	What is Tommy's favorite ice cream?	-	Chocolate	Strawberry
Test 1	Why are the Wugs sneezing?	Why are the Feps furry?	Because the Wugs are allergic to the Feps' fur.	Because the Feps go up to the mountains and the fur keeps the Feps warm.
Test 2	Why are the small Daxes wearing hats?	Why are the hat-making Blickets bigger than the Daxes?	Because the big Blickets made the hats for them.	Because the Blickets are older than the Daxes.
Test 3	How did the banana eating Gazzers learn to juggle?	How did the Gazzers get the bananas?	Because Bozo the clown taught the banana eating Gazzers how to juggle.	Because the Gazzers threw their balls up into the trees and knocked down the bananas.
Test 4	How did the Duff's toy fall into the deep hole?	How did the Duffs rescue their toy?	Because the Duffs' hair was in their eyes and they couldn't see, and they tripped and dropped their toy.	Because the Duffs tied their long hair into a rope and made a ladder with it and used it to climb down and get the toy.

Note. Participants were excluded for responding incorrectly to either Training trial.

3.2.2 Results

In Experiment 1a, our primary research question was whether participants would choose the appropriate explanation on each trial: Facts for Factual questions and Conjectures for Conjectural questions. Figure 3-2 shows children's responses by question type. Across all age groups and conditions, children successfully matched answers with question types (3.17 of 4 trials; SD = 0.71; $t(65) = 13.27$, $p < .001$). A third of the children (23/66) chose the appropriate answer at ceiling (binomial $p < .001$).

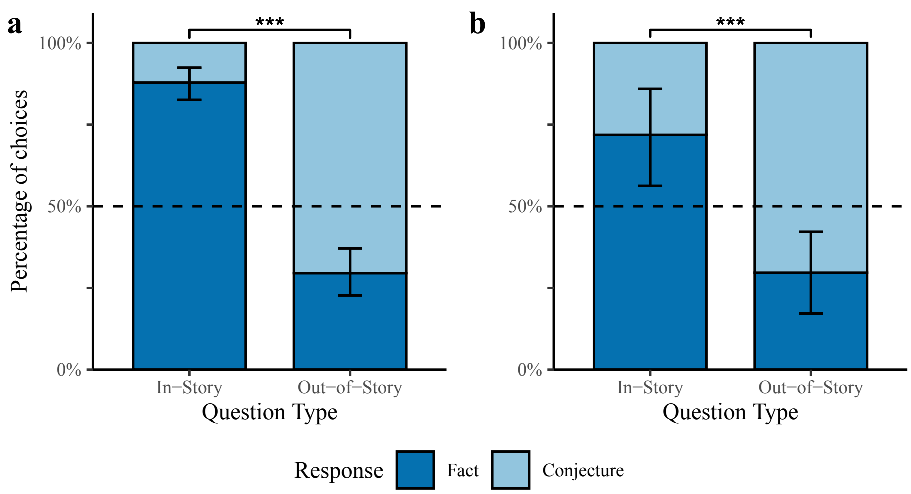

Figure 3-2: Children's ratings (averaged across two test trials for each question type) in **Experiments 1a** (N = 66, mean: 6.04 years; range: 4.00–7.93) and **1b** (N = 32; mean: 5.03 years; range: 4.15–5.92). Children were more likely to choose facts when the question could be answered by information in the story and conjectures when it could not. Error bars show bootstrapped 95% confidence intervals. Paired t-test, ***$p < .001$

We looked at whether children's responses varied by age and question type using a logistic mixed-effects model. This model predicted children's response (0=fact, 1=conjecture) from age (in months, mean-centered), question type (0=In-Story, 1=Out-of-Story), and an interaction of age and question type, with random intercepts for subject and story. There was no effect of age ($p = .080$ by asymptotic Wald test) or an interaction of age and question type ($p = .155$). As predicted, there was a main effect of question type (β=3.23; OR=25.24, 95%CI=[11.9, 53.7]; $p < .001$). Children endorsed the fact more often on In-Story questions (M = 1.76 of 2 trials, SD = 0.43) and the conjecture more often on Out-of-Story questions (M = 1.41 of 2 trials, SD = 0.58).

Given that there was no effect of age in Experiment 1a, we looked at whether the same results would hold looking only at the four and five-year-olds. Experiment 1b was identical to Experiment 1a, except that, as noted above, we presented "how" stories before the "why" stories (see Appendix A for post-hoc exploratory analyses of performance by "how" and "why"). Four and five-year-olds successfully matched answers with question types (M=2.84 out of 4 trials, SD = 1.14; $t(31) = 4.19$, $p < .001$). As in Experiment 1a, approximately one-third of children (12/32 or 38%) chose the appropriate explanation at ceiling ($p < .001$).

Next, we fit a logistic mixed-effects model to predict children's choices from age, question type, and age by question type interaction, with random intercepts for subject and story. Replicating Experiment 1a, we found a main effect of question type (β=2.22; OR=9.24, 95%CI=[3.4, 25.4]; $p < .001$); children endorsed the Fact more often on In-Story questions (M = 1.44 of 2 trials, SD = 0.84) and the Conjecture more often on Out-of-Story questions (M = 1.41 of 2 trials, SD = 0.71). There was no effect of age ($p = .281$) or age by question type interaction ($p = .181$).

3.2.3 Discussion

The results of Experiments 1a and 1b suggest that children as young as four and five flexibly consider the explanatory demands of the question under discussion. Children did not show a consistent preference either for known facts or novel information. Rather, children appropriately used the question to guide their evaluation of possible answers. When questions could be answered by available facts, children preferred factual answers; when they could not, children rejected the established facts in favor of conjectural claims for which they had no independent evidence.

3.3 Experiment 2

Experiments 1a and 1b showed that four- and five-year-olds were willing to answer questions with novel unverified conjectures rather than known facts; however, the forced choice design meant we cannot tell whether children actively endorsed conjectures for otherwise unanswered questions or whether they simply rejected facts that failed to answer questions satisfactorily. In Experiment 2 we ask children to rate each response independently and manipulate explanation type as a between-subjects comparison such that children never got to compare facts against conjectures.

Also, in the preceding experiments, we did not give children any direct information about the empirical status of the conjectures; it is possible that the children in Experiment 1 may have accepted the conjectures because they failed to recognize that they were indeed speculative and unverified. In Experiment 2, we add an Uncertain Conjectures condition where we emphasize the speculative nature of the conjectures by prefacing them with explicit uncertainty markers ("I don't know, but maybe . . . "). Abundant evidence suggests that four and five-year-olds preferentially endorse claims from speakers who are knowledgeable and confident over those from speakers who admit ignorance or uncertainty (saying "I don't know", "Hmm", or "maybe"; Jaswal & Malone, 2007; Moore et al., 1989; Sabbagh & Baldwin, 2001; Sabbagh & Shafman, 2009; Tenney et al., 2011). If in independent judgments, children appropriately endorse conjectures even when they are advanced by uncertain speakers, this would be a strong evidence that children value conjectures based simply on their ability to answer otherwise unresolved questions.

Note that if children succeed in this task, this would not be the first study to show that children sometimes prefer hesitant speakers to confident ones. Indeed, children show precisely this preference when a hesitant speaker is appropriately calibrated to her uncertainty (e.g., because she lacks epistemic access) and a confident (but ignorant) speaker is mis-calibrated (Birch, Severson, & Baimel, 2020; see also Huh, Grossmann, & Friedman, 2019). Critically however, there are a number of methodological differences between our task and previous work showing that children prefer informants who appropriately mark their uncertainty. Prior studies involved agents who did or not know specific facts (e.g., the contents of a box, the name of an object; Brosseau-Liard, Cassels, & Birch, 2014; Tenney et al., 2011). Here by contrast, informants are probed for explanations of causal events. We believe children might tolerate causal conjectures without explicit uncertainty markers precisely because children may recognize that such answers are speculative. When it is in common ground between the child and the informant that the relevant facts are not available, it might be less important that informants convey their uncertainty explicitly. Thus, consistent with the calibration literature, we expect that children will be un-swayed by confident statements that fail to answer a question and that children will endorse conjectures that answer questions when the informant expresses uncertainty. However, insofar as children recognize conjectures as such, we predict that they will not penalize informants who fail to convey uncertainty.

We treated this as a confirmatory study and preregistered all analyses and predictions on the Open Science Framework (`osf.io/zpq3r`). Power analysis using simulations from pilot data indicated that a sample of 32 participants per condition would yield 80% power to detect a moderate interaction of Explanation type by Question type.

3.3.1 Methods

Participants

Participants were 95 four- and five-year-olds (M = 5.03, range: 4.03-5.99) recruited and tested as in Experiment 1. Thirty-four additional children participated but were excluded for responding inaccurately on an inclusion trial (N = 26), not speaking English as a native language (N = 2), experimenter error (N = 5) or failing to complete the study (N = 1). Note that the exclusion rate of (21% of 121 initial participants) is high and similar to Experiment 1b. In this case, excluded participants overwhelmingly (21 of 26 children) correctly put Elmo in the "good cup" on the first trial and then also (incorrectly) put Cookie Monster in the "good cup". That is, in a forced choice of Elmo or Cookie Monster (Experiment 1b) children showed an "Elmo bias"; in independent judgments (and perhaps unsurprisingly given the status of these characters in children's lives) children showed a positive response bias to both puppets. Participants were randomly assigned to the Fact, Conjecture, or Uncertain Conjecture condition and we found no condition differences in age (M = 5.07, 4.98, and 5.02 years respectively; $p > .8$).

Materials and Procedure

The Materials and Procedure were identical to those in Experiment 1 except as follows. Instead of using the same pair of puppets on every trial, six different cartoon characters were used, one for each trial. The characters were taken from the Muppets and each character was printed on laminated paper and glued to a wooden stick; puppets were 15 cm tall. Each puppet appeared just once so that children could evaluate each question and answer pair independently across trials. We also used two identical blue cups (20 tall), one labeled with a smiley face sticker (in which the child could put puppets who gave "good answers") and one left blank (for placing puppets who gave "not so good answers"). These were kept in a fixed position with the "good answer" cup on the child's left and "not so good answer" cup on the child's right. See Figure 3-3.

Figure 3-3: Puppets and rating cups used in Experiments 2–4

Children were told that they would hear some stories and then hear some questions, and that the puppets would try to answer those questions. The experimenter explained, "Some puppets will give good answers and some puppets will give not so good answers". Children were introduced to the cup for good answers and the cup for not so good answers and asked to point to each. All participants correctly identified the two cups before proceeding to the training trials.

On each trial, the experimenter narrated a story accompanied by an animated slide deck and posed a question at the end. The children were introduced to just one puppet on each trial and the puppet responded with a pre-recorded answer (activated by the experimenter). The experimenter then asked the child, "Was that a good or not so good answer?" Children rated the puppet's response by placing them into one of two cups.

As in Experiment 1, we designed the training trials to familiarize participants with the question and explanation evaluation process, and to elicit both ratings. Training trials were the same for every child: the puppet on the first trial provided a good answer to the question; the puppet on the second trial provided a not so good answer. Children received feedback on these items to reinforce the two-cup rating system. Any child who responded incorrectly on either training trial was excluded from further analysis and replaced. Next, the experimenter presented the four test trials. As in Experiment 1, two trials involved In-Story questions and two trials involved Out-of-Story questions.

In the Fact condition, the puppet on each of the four trials responded with a verified true fact from the story (i.e., regardless of whether they were asked an In-Story or Out-of-Story question). Thus, on the In-Story trials, the puppets provided good answers and on the Out-of-Story trials, the puppets provided not so good answers. In both the Conjecture and Uncertain Conjecture condition, the puppets on each of the four trials responded with an unverified conjecture; thus, in both conditions, the puppets on the In-Story trials provided not so good answers and the puppets on the Out-of-Story trials provided good answers. In the Uncertain Conjecture condition, the answer was preceded by "I don't know, but maybe ...". Test trials were counterbalanced as in Experiment 1, with eight versions per condition.

3.3.2 Results

Participants' responses are shown in Figure 3-4. As was evident in the inclusion trials children showed a positivity bias towards both puppets, children rarely placed any of the puppets in the "not so good" cup. Across all conditions, children successfully matched answers with question types (M = 2.38 of 4 trials; SD = 0.92; t(94)=3.99, $p < .001$), although at a lower rate than in Experiment 1. About 12% of the children (11/95) chose the appropriate answer at ceiling, not significantly different than chance (binomial $p = .051$).

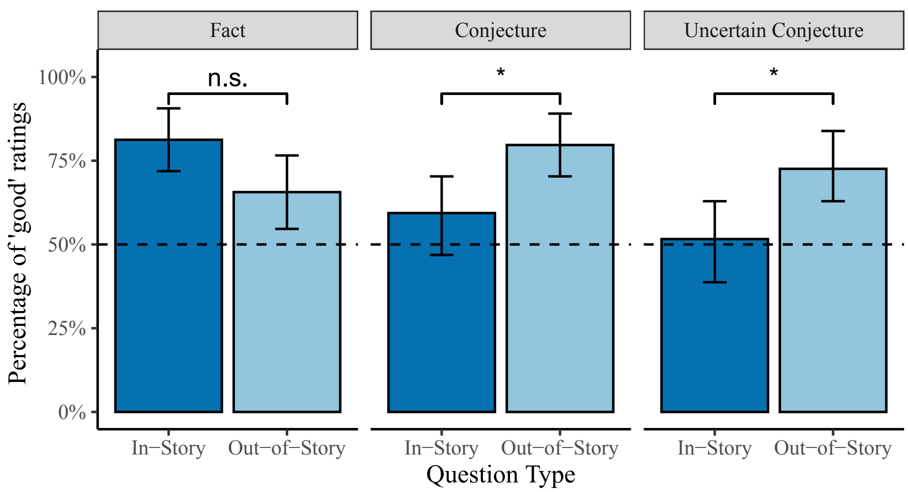

Figure 3-4: Children's (N = 95; mean: 5.03 years; range: 4.03–5.99) ratings in Experiment 2 averaged across two test trials for each question type. When given only factual answers, children tended to endorse these answers across the board. In contrast, children were more likely to endorse conjectural answers (whether offered neutrally or with an explicit uncertainty marker) only when the question could not be answered by information in the story. Error bars show bootstrapped 95% confidence intervals. (Paired t-tests, $*p < .05$)

Our first question was whether children would give Facts higher ratings for In-Story questions and Conjectures higher ratings for Out-of-Story questions. We predicted that the likelihood of endorsing each explanation would depend on an interaction between Explanation Condition (reference category=Fact) and Question Type (reference category = In-Story). To test this prediction we used a mixed effects logistic regression to predict children's endorsement on each trial, including fixed effects of Explanation Condition, Question Type, and their interaction, as well as random intercepts for subject and story (model syntax: RatedAsGood   Explanation Condition * Question Type + (1|Subject) + (1|Story)).

As predicted, the Explanation Condition by Question Type interaction explained significant variance ($\chi^2(2)= 14.9$; $p < .001$). We conducted follow-up contrasts using estimated marginal means, with Bonferroni corrections for multiple comparisons. These contrasts found that children in the Fact condition were more likely to give positive ratings on In-Story questions (84%) than Out-of-Story questions (68%), although the result did not reach the threshold for statistical significance ($p = .18$). In contrast, children hearing Conjectures were more likely to give positive ratings on the Out-of-Story questions (83%) than the In-Story questions (61%; $p = .04$). Children hearing Uncertain Conjectures were also more likely to give positive ratings on Out-of-Story questions (76%) than the In-Story questions (52%; $p = .05$).

Our second question was whether children's ratings of conjectures would be affected by expressions of uncertainty. Follow-up contrasts comparing children's ratings in the Conjecture and Uncertain Conjecture conditions found no difference for either In-Story questions or Out-of-Story questions (z's < 1). We also tested for any interaction of question type and condition type by repeating the previous regression analysis including just the Conjecture and Uncertain Conjecture conditions. This regression analysis did not find a significant Explanation Condition by Question Type interaction (z = -0.06, $p = .954$) or a main effect of Explanation Condition (z = -0.83, $p = .404$). However, there was a main effect of Question Type ($\beta = 1.23$, OR = 3.43, 95%CI = [1.40, 8.41], z = 2.69, $p = .007$), showing that children consistently rated conjectures more highly for Out-of-Story questions than In-Story questions. Thus, children's judgments were not significantly impacted by explicit expressions of uncertainty.

finally, we asked whether children's sensitivity to the match between question type and explanation type was driven by an active recognition of appropriate answers, a rejection of inappropriate answers, or both. We used one-sided Wilcoxon signed rank tests to compare children's ratings against chance for each of the 6 combinations of Explanation Condition by Question Type (i.e. each bar in Figure 3-4). Correcting for multiple comparisons, we found that children rated explanations "good" significantly more often than chance when explanations were appropriate to the question type (e.g. Facts for In-Story questions and Conjectures for out-of-Story questions, p's $<$ .0083). However, when explanations were inappropriate to the question type, children did not reject these explanations more often than chance (p's $>$ .0083).

3.3.3 Discussion

Although children in Experiment 2 were inclined to endorse all the answers they were given, they were nonetheless sensitive to the relationship between questions and answer types. Replicating Experiment 1, children in Experiment 2 preferentially endorsed facts for questions whose answer could be found in the story and conjectures for questions whose answer was unknown. Critically in this context, adding explicit markers of ignorance and uncertainty ("I don't know, but maybe ...") did not impact children's endorsement of conjectural explanations.

These results are compatible with a growing literature showing that children use situational constraints to evaluate testimony; that is, children assess whether speakers' claims are justified given their epistemic access (e.g., Birch et al., 2020; Brosseau-Liard et al., 2014; Huh et al., 2019; Tenney et al., 2011; see Koenig, Tiberius, & Hamlin, 2019, for review. Consistent with this literature, we find that four- and five-year-olds readily accept hesitant speakers who offer speculations as answers to unresolved questions. However, in contrast to previous work suggesting that children penalize overconfident speakers who fail to convey hesitancy when reporting uncertain information, children here were happy to endorse speakers who advanced conjectures without uncertainty markers. As noted, we believe that this is because the uncertainty marker might be redundant in these contexts insofar as both the children and informant recognize that the answers were indeed speculative.

One limitation in interpreting the results of Experiment 2 is that children were generally inclined to endorse all the answers that they were given. Future research might use a more sensitive measure of children's judgments, such as a rating scale with more than two response options, or by asking children to explain their ratings. For consistency and ease of comparison however, we used the same binary rating scale in Experiments 3-4.

3.4 Experiment 3

In Experiment 2, children endorsed conjectures that contained information not substantiated by the stories. However, although the information was novel, it was not especially surprising. Would children be willing to endorse conjectures that answered otherwise unanswered questions if the conjectures were improbable given the children's prior expectations?

Abundant research has testified to young children's ability to integrate evidence with prior knowledge to draw rational inferences (Gopnik & Wellman, 2012; C. Legare, Gelman, & Wellman, 2010; Schulz, 2012b; D. Sobel & Kushnir, 2013; Tenenbaum, Kemp, Griffiths, & Goodman, 2011; Xu & Tenenbaum, 2007). Children's sensitivity to both the prior probability of hypotheses and the data in their favor might reasonably lead children to reject conjectures that are supported by neither.

However, a remarkable feature of human learning is that we can and do go beyond both the data and current knowledge to advance new, and even initially prima facie unlikely ideas. After the explosion on Apollo 13, the astronauts and ground crew had to improvise a way of connecting canisters for removing carbon dioxide to the lunar module; they succeeded using cardboard ripped from their training manual, towels and duct tape. The proposal was endorsed, not because of the weight of evidence in its favor, nor because cardboard, towels and duct tapes were typically used for these ends, but because – faced with a problem and no apparent solution – a speculative proposal that might solve the problem could be valued on those grounds alone. Obviously, Apollo 13 was an extraordinary event; in everyday cognition, our willingness to endorse otherwise unfounded conjectures may be more likely to inflame superstitions or perpetuate conspiracy theories than save lives. Still, the ability to value hypotheses simply because they could answer questions or solve problems (an ability that we, in homage to William James, will refer to as cognitive pragmatism) may contribute to human learners' distinctively powerful ability to generate new knowledge about the world.

In Experiment 3, we look at participants' willingness to accept conjectures that answer the question at hand when the conjectures contradict expectations set up by the story or when the conjectures are also rare, low probability events in themselves. That is, rather than pitting conjectural answers against facts, we compare more and less plausible conjectures, in contexts where they either do or do not provide answers to the question at hand. We test both adults and a relatively wide age-range of children (as in Experiment 1a), since to our knowledge, no studies have looked at whether participants at any age will prioritize the pragmatic goal of answering questions over other considerations and endorse unlikely conjectures when they provide a potential resolution to otherwise unresolved queries. An a priori power analysis was conducted using G*Power3 to test main and interaction effects in a 2x2 repeated measures ANOVA, using an F-test with a medium effect size ($f = .25$) and an alpha of .05. Results indicated that a sample of 24 participants was required to achieve a power of 0.8.

3.4.1 Methods

Participants

Twenty-four adults were recruited and tested via Amazon Mechanical Turk and paid $1.00 for participating. Ten additional adults were excluded for failing to complete the experiment ($n = 2$) or failing to distinguish good and bad responses on the two inclusion trials by at least a ten-point spread ($n = 8$). Twenty-four four- to seven-year-olds ($M = 5.95$ years, range = 4.37-6.71) were recruited and tested as in the preceding experiments. Fourteen additional children participated but were excluded for responding inaccurately on the two inclusion trials ($N = 7$), being distracted during one or more trials ($N = 1$) or failing to complete the study ($N = 6$).

Materials

We created six test stories, each paired with one question and four candidate answers. Twenty-four unique puppets were used to present the candidate answers, with characters from the Muppets and Hey Arthur! shows and constructed like those in Experiment 2. The two “good answer” and “not so good answer” cups used in Experiment 2 were also used here.

In contrast to the previous studies where we wanted to control for prior knowledge, in Experiment 3 we wanted to leverage participants’ background knowledge to evoke strong expectations about plausibility. Thus, rather than using novel characters, these stories involved human children in everyday activities. The stories always ended with the protagonist in a salient emotional or behavioral state and the questions all asked why the character was in that state. The story context always set up one answer as the most plausible answer (i.e. the “Likely Answer”). However, in contrast to the previous studies, no question was directly answered by the information in the story (i.e. we only asked Out-of-Story questions). All the candidate explanations were conjectures; they varied in whether or not they answered the question at hand and how plausible they were. That is, on every trial, the four candidate answers always crossed two factors: whether it would answer the question if it were true (Answer / Non-answer) and how probable the conjectured event was (Likely / Unlikely). To ensure that children recognized which candidate answers violated expectations, all answers (except for the Likely Answer) began by explicitly denying the likely answer (i.e. they took the form "Not because of ... but because of ..."). See Table 3.2 for the complete text of one trial; full stimuli are presented in Appendix A.

As in preceding experiments, we included two training trials to familiarize participants with evaluating the answers to questions about the stories. Although we had used only fact questions in the training for Experiments 1 and 2; in the training for Experiment 3 we included a conjectural question because all of the test trials involved contrasts between conjectures. On the first training trial, we asked a factual question that could be answered by recalling information from the story. Participants rated two answers: one was true with respect to story and answered the question; the other was true but did not answer the question. On the second training trial, we asked a conjectural question that could not be answered given available information. Participants rated two explanations: one was a likely answer given the story; the other was unlikely. (Note therefore that if anything, the training trial should make participants less likely to endorse unlikely conjectures.) These training trials were the same for all participants and were used as inclusion criteria. To ensure that adults on M-Turk were following the task instructions, we used a fairly conservative inclusion criteria: adults had to distinguish the good and bad questions by a 10 point spread in ratings; adult participants were excluded from further analysis if they failed to make this distinction (and were thus possibly responding at random) or if they reversed the ratings. children were excluded from further analysis and replaced if they answered either training trial incorrectly.

Procedure

We used a within-participants design: participants saw all six test stories and rated all four answers for each story. The stories were presented in a fixed sequence; answers were randomized for adults and presented in pseudorandom order for children.

Adult participants used a linear rating scale ranging from Not Satisfying (0) to

Table 3.2: The two training trials and one of six test trials in Experiment 3.

Item	**Text**
Training Story 1	This is Tina. Tina is having breakfast. She's eating pancakes. After breakfast, she went outside to play. When she was done, it was time for lunch. Tina comes back into the kitchen. Her brother comes into the room and asks, "Hey Tina, what did you have for breakfast?"
Question	What did Tina have for breakfast?
Likely Answer (Fact)	Tina ate some pancakes for breakfast.
Likely Non-Answer	Tina played in the tree-house for breakfast.
Training Story 2	This is Tommy. Tommy went to the ice cream shop. He bought his favorite ice cream and ate it all up! (only an empty cone was shown)
Question	What did Tommy get at the ice cream shop?
Unlikely Non-Answer	Tommy got tomato soup.
Likely Answer	Tommy got chocolate ice cream.
Test Story	This is Sally. Sally was looking forward to her best friend's birthday party. Her best friend had just mailed out the invitations, and Sally was hoping to get one soon. Sally walked to her mailbox and saw a shiny white envelope. Sally opened the envelope and jumped up and down excitedly when she read it!
Question	Why was Sally so excited?
Likely Answer	Because she got invited to her best friend's birthday party. She was so excited that she couldn't stop jumping up and down.
Unlikely Answer	Not because it was a birthday invitation, but because it was a letter from school saying she won the story competition.
Likely Non-Answer	Not because it was a birthday invitation, but because it was a notice from the library saying she forgot to return her books.
Unlikely Non-Answer	Not because it was a birthday invitation, but because it was a note from her teacher saying that she had to do extra work after school.

Note. See Appendix A for full stimuli. Participants were excluded for incorrectly responding to the training trials. In Experiment 4, we used the same training trials and test stories and questions, but only presented the Unlikely Answers and the two Non-Answers. We also modified the conjectures in Experiment 4 to the form "It was not [a birthday invitation /. . .]. Sally was excited because [it was a letter from school saying she won the story competition /. . .].".

Very Satisfying (100) and saw all four conjectures presented on the same page. Adults did not receive any feedback on the training trials but adults were excluded from further analysis and replaced if they did not rate the better explanation at least 10 points higher than the alternative on the training trials.

Children were tested individually in a quiet room with a laptop computer. The experimenter provided feedback on training trials to reinforce the two-cup rating system. Any child who responded inappropriately on a training trial was excluded from further analysis and replaced. The procedure and rating system was identical to the one used in Experiment 2 except as follows. The experimenter read each story out loud and then asked the target question. Four puppets took turns giving each of the candidate responses (as in the previous studies, all responses were pre-recorded to avoid differential prosodic cues across trials). Children rated each conjecture as either a "good" or a "not so good" answer by placing the puppet into the appropriate cup. Children had to rate each puppet's answer before hearing the next puppet's answer.

3.4.2 Results

Our main question was whether participants would evaluate conjectures based not just on how plausible the conjecture was but also on how well it would answer the question, if true. We used mixed effects regression predicting responses to each conjecture. This analysis included fixed effects of Answer Type (Answer / Non-Answer) and Likelihood (Likely / Unlikely), as well as random effects for subject and story (Model syntax: Response Answer Type * Likelihood + (1|SubjectID) + (1|Story)). Recall that adult participants provided a continuous rating (0-100) but child participants provided a binary rating (0 or 1). Thus, we used linear regressions for adults and logistic regressions for children.

Among adult participants, there was a significant effect of Answer Type (β = 33.41, 95% CI=[27.9, 38.9], z=11.89, $p < .001$), with higher ratings for Answers (M=64.9, SD=14.4) than Non Answers (M=10.6, SD=19.8). The effect of Likelihood was not significant (z=0.05, $p = .961$), however, there was a significant Answer Type by Likelihood interaction (β = 41.73, 95% CI=[33.9, 49.5]; z=10.50, $p < .001$).

We inspected the interaction using follow-up Tukey contrasts correcting for multiple comparisons. Although adults gave higher ratings to Likely Answers (M=85.81, SD=18.37) than Unlikely Answers (M=43.94, SD=26.02; $p < .001$), they did not differentiate between Non-Answers that were Likely (M=10.7, SD=18.4) or Unlikely

(M=10.5, SD=21.8; $p > .99$). Critically, adults endorsed Unlikely Answers more often than Unlikely Non-Answers ($p < .001$) and also Likely Non-Answers ($p < .001$). See Figure 3-5a for adults' average ratings.

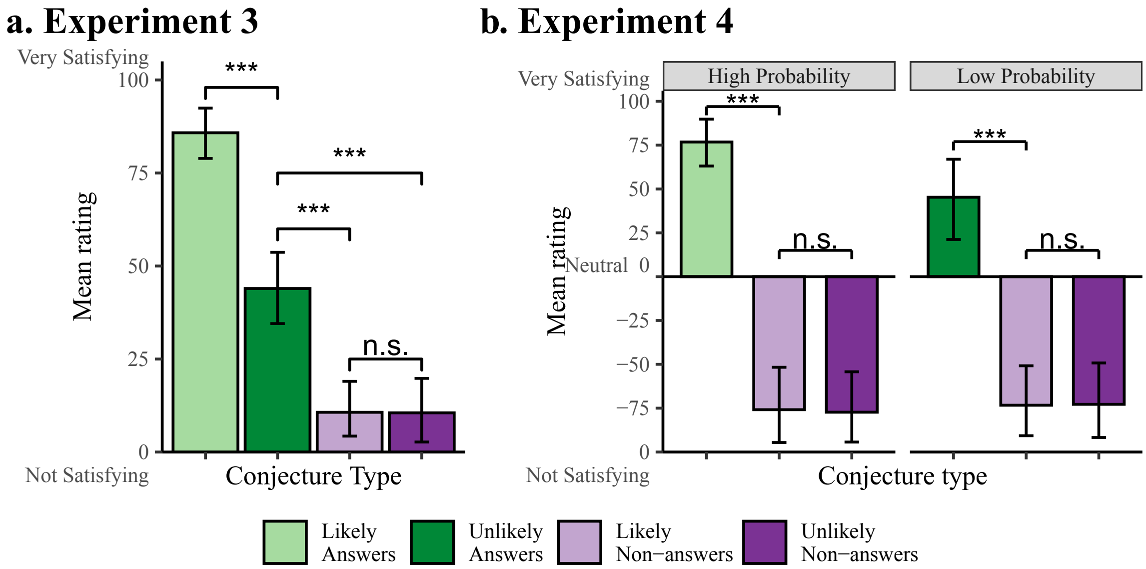

Figure 3-5: Adults' ratings in (a) Experiment 3 (N = 24) and (b) Experiment 4 (N = 24) averaged across the six test trials. When given all responses and a positive rating scale (0-100), adults preferred likely answers to unlikely answers and both of these to non-answers (Experiment 3). The results replicated when given non-answers and only either likely or unlikely answers, and a scale allowing answers to rate answers either negatively or positively (-100 to 100, Experiment 4). Participants rated unlikely answers nearly as positively as likely ones. Error bars show bootstrapped 95% confidence intervals. (Tukey's HSD, *** $p < .001$)

The results for children were similar (see Figure 3-6). As in the adult sample, there was significant effect of Answer Type ($\beta = 1.38$, OR=3.96, 95% CI=[1.81, 8.65], z= 3.45, $p < .001$), no main effect of Likelihood (z=.22, $p = .82$), and a significant Answer Type by Likelihood interaction ($\beta = 3.68$, OR=39.53, 95%CI=[12.28, 127.21]; z = 6.17, $p < .001$). Follow-up Tukey contrasts using estimated marginal means indicated that children were more likely to endorse Likely Answers (M=90%, SE=4%) than Unlikely Answers (M=17%, SE=5%; $p < .001$), but did not differentiate between the Likely Non-Answers (M=5%, SE=2%) and Unlikely Non-answers (M=5%, SE=2%; $p > .99$). Children also endorsed Unlikely Answers more often than both Likely Non-Answers ($p = .049$) and Unlikely Non-Answers ($p = .040$).

We also looked at the effect of age on children's responses by adding a covariate of age (in years, mean-centered) to the logistic mixed effects model predicting children's

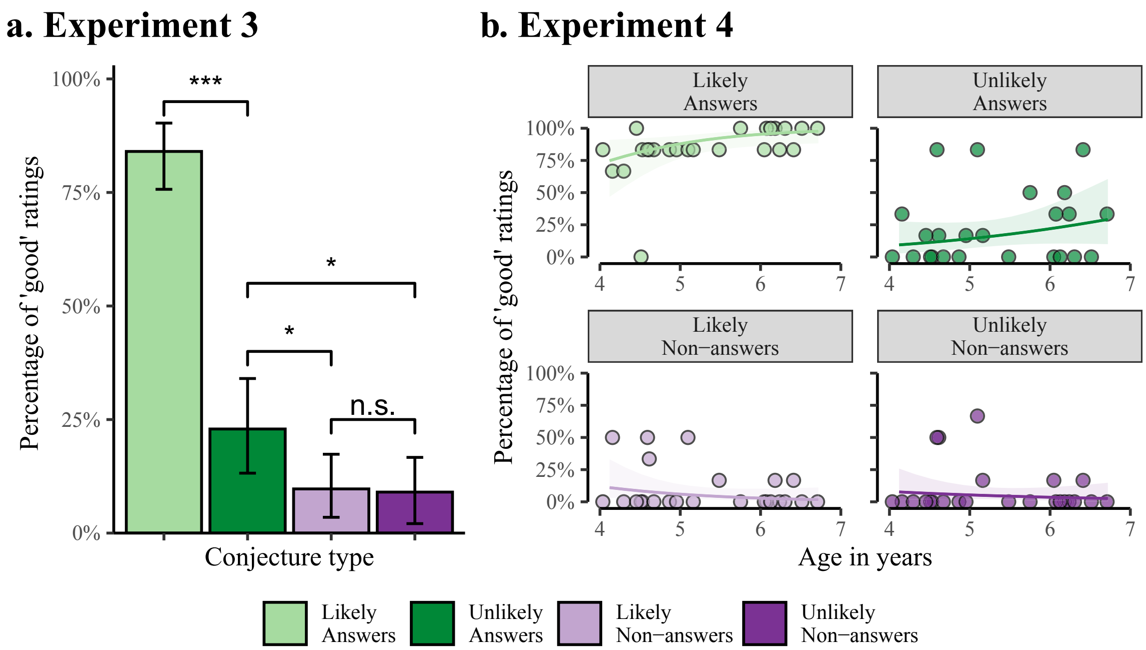

Figure 3-6: Percentage of conjectures rated as good by children in Experiment 3 (N = 24; mean: 5.95, range 4.37– 6.71), averaged across the six test trials. Children rated all four conjecture types on each trial. (a) Responses averaged across all participants; (b) The same data by answer type and child's age; each circle represents responses from one child on that answer type. Like adults, children preferred likely answers to unlikely answers and both of these to non-answers. With increasing age, children were more likely to accept Answers and more likely to reject Non-Answers, for both likely and unlikely conjectures. The distinction between Likely and Unlikely Answers decreased with age. Lines show predictions from regression model; shaded regions and error bars show 95% confidence intervals. (Tukey's HSD, *$p < .05$, ***$p < .001$)

responses and including an age by conjecture interaction. Because we did not have specific hypotheses about how age might interact with children's ratings for each conjecture type, we used a categorical variable of Conjecture Type (with Likely Answers as the reference level) so that estimated coefficients could be directly interpreted as the effect of age on the log odds of endorsing each conjecture type. The model thus included fixed effects of conjecture type, age, and a conjecture type by age interaction, with random intercepts for subject and story. This expanded model explained significant additional variance than the original analysis without age (χ^2 (4)= 14.36, p = .006 by likelihood ratio test), although the overall effect of Conjecture type remained significant ($\chi^2$2 (3)=124, $p < .001$). Critically, this is qualified by an age by conjecture type interaction ($\chi^2$2 (3)=13.53, p = .004). Inspection of estimated marginal slopes indicated that with increasing age, children were more willing to endorse conjectures that answered the question, regardless of how otherwise plausible they were (Likely Answers: β= 1.02, 95%CI=[.08, 1.96]; Unlikely Answers: β = .53, 95%CI=[-.30, 1.37]) and less likely to endorse conjectures that did not answer the question, again regardless of how otherwise plausible they were (Likely Non-Answers: β = -.77, 95% CI=[-1.86, .32]; Unlikely Answers: β = -.43, 95%CI=[-1.49, 0.62]). Figure 3-5b illustrates these marginal effects of age on the probability of endorsing each conjecture type.

3.4.3 Discussion

The results of Experiment 3 suggest that both adults and four- to six-year-old children will endorse even otherwise unlikely conjectures as long as they offer potential answers to questions. This is not to say that people are indifferent to the plausibility of conjectures: Both adults and children preferred conjectures involving likely events to those involving unlikely ones. Importantly however, and consistent with the idea of cognitive pragmatism, participants' evaluations privileged the degree to which a conjecture might answer the question, and only secondarily considered how likely the conjecture might be.

However, although adults rated the unlikely answers as more satisfying than the non-answers, the average rating of unlikely answers was close to the middle of the scale, suggesting that adults might not so much have endorsed as merely been indifferent to the unlikely conjecture. In Experiment 4, we replicate the design of Experiment 3 but use a -100 to +100 rating scale to allow us to distinguish adults' active endorsement of unlikely conjectures from more neutral or negative responses. We also gave both adults and children just one kind of answer, plausible or implausible, to see how participants might evaluate prima facie implausible conjectures when they are not explicitly contrasted with more plausible ones.

3.5 Experiment 4

In Experiment 4, we used a between-participant design for adults and a within-participant design for children. We asked both adults and children to evaluate three conjectures for each question: one low probability conjecture that answered the question (Unlikely Answer) and two conjectures that did not answer the question (i.e. Likely and Unlikely Non-Answers). In adults, we also ran a condition in which participants rated a high probability conjecture that answered the question (Likely Answer) and both Non-Answers. For adults, our primary question of interest was whether there would be any difference in their ratings for Likely versus Unlikely Answers between conditions. For children, our primary aim was to follow-up on the age effect and see whether older children would be more likely than younger children to endorse low-probability conjectures that answered the questions at hand (Unlikely Answers). Based on the large effect size in Experiment 3 (Cohen's d = 1.73), we aimed to recruit the same number of participants per sample for this follow-up. Specifically, 12 adults per condition would provide 98% power to detect the same effect size and 80% power to detect a smaller but still substantial effect of d = 1.2. For children we aimed to recruit 12 four and five-year-old and 12 six-and seven-year-olds to yield 80% power to detect an age by conjecture type interaction with a large effect size (from Experiment 3: partial $\eta^2 = 0.25$; Cohen's f = 0.58).

3.5.1 Methods

Participants

Twenty-four adults were recruited and tested via Amazon Mechanical Turk and paid $1.00 for participating. An additional thirteen adults were excluded for failing to distinguish good and bad responses on the inclusion items by at least a ten-point spread (n = 12) or failing attention checks (n = 1). Twenty-four four- to eight-year-olds (M = 5.32 years, SD = .86, range = 4.04 – 6.71) were recruited and tested as in the preceding experiments. Six additional children participated but were excluded for responding inaccurately on the inclusion questions (N = 5) or failing to complete the study (N = 1).

Materials and procedure

The materials and procedure were the same as in Experiment 3 with three modifications. first, in Experiment 4, participants rated three conjectures per test trial instead of four. Thus, for children, only three puppets were used per trial, for a total of 18 puppets. Second, we modified the rating scale for adult participants to range from -100 ("Extremely Dissatisfying") to +100 ("Extremely Satisfying), with 0 explicitly marked as "Neutral". Children used the same binary response measure from Experiment 2-3, rating each puppet's proposal as a good or not so good answer. Third, we modified the sentence structure of the conjectures. In Experiment 3 we had used one long sentence of the form "Not because [likely answer], but because [conjecture]". In Experiment 4, we separated the denial and the conjecture into two sentences (e.g., "Sally was not excited because it was a birthday invitation. Sally was excited because it was a letter from school saying she had won a story competition.") both to remind participants of the target of the explanation and to make the denial of the likely explanation even more explicit.

3.5.2 Results and Discussion

We had two primary questions. first, we wanted to know whether adults would distinguish likely and unlikely conjectures in their evaluations if each conjecture were presented independently. Adults gave higher ratings to Likely Answers (M=76.8, SD=15.4) than Unlikely Answers (M=45.3, SD=43.6; t(17.7) = 2.16, p = .04). Importantly, although participants recognized that one answer was more plausible than the other, they treated unlikely conjectures as acceptable answers when the improbable conjecture was presented as the only answer to a question: of the 12 adults rating Unlikely Answers, 10 gave positive ratings, and 3 gave ratings above 90 (comparable to the 5 of 12 adults who gave ratings above 90 in the Likely Answer condition). Consistent with the idea of cognitive pragmatism, adults seemed to value the degree to which a proposal could answer a question above the prior probability of the proposal. See Figure 3-5b for adults' average responses.

Like adults, children in Experiment 4 endorsed Unlikely Answers (M = 44%) more often than both Likely Non-Answers (M=15%, Wilcoxon signed rank test $p < .001$) and Unlikely Non-Answers (M=13%, $p < .001$). Again however, this was qualified by an age by conjecture type interaction. We built a logistic regression model with the fixed effects of conjecture type, age, and an age by conjecture type interaction. As planned, we analyzed age as a binary variable (ages 4-5 vs. 6-7), although similar results obtain with age coded continuously (see Appendix A for details). This model explained more variance than either the conjecture only model ($\chi^2(3) = 21.4$, $p < .001$) or the model with conjecture and age but no interaction ($\chi^2(2) = 21.3$, $p < .001$). Inspection of estimated marginal slopes indicated that older children were more likely than younger children to endorse Unlikely Answers (β=.34, p = .02), but equally likely to reject conjectures that did not answer the question (Likely Non-Answers: β= -.091, p = .18); Unlikely Non-Answers: β= .088, p = .16). Thus, and in contrast to the idea that younger children might be if anything more drawn to novel and speculative responses than older children, four and five-year-olds were in fact more likely to endorse answers that were probable in themselves whereas the six and seven-year-olds were able to represent and evaluate proposals based only on the abstract fit between the question and the answer. Figure 3-7 shows the distribution of responses by age group.

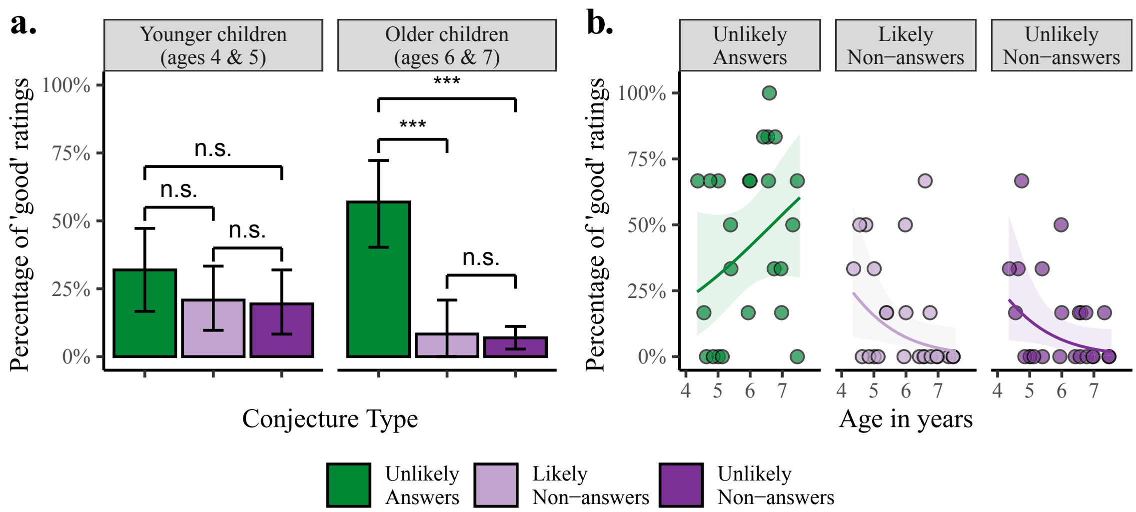

Figure 3-7: Percentage of conjectures rated as good by children in Experiment 4 (N = 24; mean: 5.32; range 4.04– 6.71), averaged across the six test trials. Children rated all three conjecture types on each trial. (a) Responses averaged across participants by age group. (b) Each circle represents responses from one child on that conjecture type. With increasing age, children were more likely to accept Unlikely Answers and less likely to accept both Likely and Unlikely Non-Answers. Lines show predictions from regression model; shaded regions and error bars show 95% confidence intervals. (Wilcoxon signed rank test, $***p < .001$)

Finally, we note that in both Experiments 3 and 4, four and five-year-olds correctly endorsed likely answers, correctly rejected both likely and unlikely non-answers, and (unlike adults and older children) sometimes also rejected unlikely answers. By contrast, in Experiment 2, four and five-year-olds tended to rate all responses (facts and conjectures) positively, although they correctly endorsed facts more often for questions that could be answered by those facts in the story and conjectures when they could not. Overall however this pattern of results raises the question of why four and five-year-olds did sometimes reject responses in Experiments 3 and 4 but generally (if differentially) endorsed all responses in Experiment 2. We believe two factors might have contributed. first, Experiment 2 never involved any conflict with the children's prior knowledge; this might have made it easier for children to endorse responses across the board. Second, all the responses except the likely response were prefaced with an explicit rejection of at the likely response ("It wasn't a birthday invitation. Sally was excited because …"). Since the experimenter effectively modeled rejecting a candidate response, the children might have felt more licensed to do so as well.

3.6 General Discussion

Across four experiments, we found that both adults and children were willing to entertain novel and unverified claims when (and only when) they answered the question at hand. Four- and five-year-olds endorsed such conjectures in both forced choice (Experiment 1) and independent judgment paradigms (Experiment 2), even when the conjectures were accompanied by explicit expressions of uncertainty (Experiment 2). Adults and six- and seven-year-old children, but not younger children, further endorsed conjectures that were improbable but provided potential answers to the questions at hand (Experiments 3 & 4). When confronting otherwise unanswered questions, adults and young children judged speculative conjectures not on the basis of evidence for their truth, but instead, on their potential for addressing those unanswered questions if they were true.

Our results add an important perspective to the growing literature on trust and testimony, which has largely found that children rationally integrate multiple cues to speaker and information reliability when deciding whether to accept or reject novel claims (for review, see Harris et al., 2018; Koenig et al., 2019; Stephens, Suarez, & Koenig, 2015). In contrast, here we find that participants endorsed novel conjectures despite conflicts with cues to reliability, including verifiability, speaker confidence, and consistency with prior expectations. Our results suggest that the primary driver of learners' judgments was whether the proposed conjecture answered the question at hand. Indeed, in the absence of other potential answers to a question, adults judged implausible conjectures just as favorably as plausible conjectures.

We also found evidence for a developmental trajectory in which older children and adults were more willing to accept low probability conjectures than younger children (Experiments 3-4). Importantly, this was the case only when the low probability conjectures answered the question; participants tended to reject all non-answers. There are at least two mutually compatible explanations for this finding. First, there may be an age-related increase in participants' tolerance for uncertainty and the value they give to getting a question answered - in other words, a move from an empirical stance to a pragmatic stance. This hypothesis is supported by young children's success at recognizing conjectures that answer a question in Experiments 1 and 2. In fact, this developmental finding is compatible with other reports of an age-related increase in the subtlety and sophistication with which children evaluate claims and testimony (for review, see C. Mills, 2013). Whereas 3-year-olds reject sources with any evidence of historical inaccuracy, 4-year-olds differentiate between sources with 75% vs. 25% inaccuracy and preferentially trust the more reliable (Pasquini et al., 2007). Thus, with age, children may learn to balance multiple criteria for deciding which claims to endorse or reject. Future work may explore how children integrate various criteria – empirical, social, and structural – when evaluating novel claims and how properties of the specific question under discussion influences which criteria takes priority.

Second, there may be an age-related improvement in participants' ability to recognize claims that answer a question. Precisely how people assess the satisfactoriness of a conjecture in the absence of any evidence remains an open question; we speculate that people might be broadly sensitive to information contained in the question itself. For instance, we know that by age three, children recognize that interrogative words specify desires for different categories of information (Ervin-Tripp, 1970), such that "who" requires an agent and "where" requires a place; Piaget (1926) further divided "why" questions into subtypes: causal explanation, human motivation, justification, and logical explanations. More generally, beyond linguistic information, the ability to evaluate novel conjectures might be related to developing world knowledge (e.g., in Experiment 3, an understanding of what events might make someone excited or sad) or problem-solving abilities (e.g., in Experiment 1, knowing that you could retrieve out-of-reach objects by knocking them off the platform with a ball). Over time, accumulated linguistic experience, world knowledge, and planning abilities might help children fine-tune their understanding of what different questions require of their answers.

We have contrasted the current work with past work on trust in testimony. Here we show that so long as a response provides a potential answer to a question, children are willing to override known fact (even when communicated confidently) in favor of speculative conjectures (even when advanced uncertainly). However, we do not mean to suggest any fundamental incompatibility between our results and the literature on trust in testimony. Children might well be epistemically vigilant in tracking informants who do and do not provide satisfactory answers to questions, even when all the answers are conjectural. In future research it might be interesting to see if children track agents' history of answering questions with appropriate (versus irrelevant) conjectures and whether children use this information to make decisions about whom to address when posing questions likely to have no known answer.

It should be noted that although across these studies we tested a variety of scenarios (made-up monsters and real-life events that differed in plausibility) and used a range response measures (forced choice and independent rating), we lack direct evidence that children in Experiments 3 and 4 explicitly understood the speculative nature of conjectures. Given that young children can be both overly credulous and overly skeptical of unfamiliar events (Woolley & Ghossainy, 2013), more research is needed to examine children's credulity towards speculative claims. For example, in Experiment 3 we might ask children to explain their ratings and explicitly judge the probability of the conjectures ("Do you think that really happened?") or the speaker's belief in the conjecture ("Does the speaker believe that it really happened?").

We began this line of work by observing that remarkable discoveries often emerge from wild speculations in response to novel, unanswered questions. In science, religion, and everyday life, we give license to conjectures, speculations, and other unverified assumptions when they can play powerful explanatory roles. An important next step is to go beyond looking at how people evaluate claims to looking at how they generate them. Any number of factors might affect people's willingness to engage in speculative reasoning. For example, people might be more willing to generate conjectures in response to scientific questions than religious ones if they believe that accepting mysteries in religion is a sacred value; on the other hand, they might be more willing to generate conjectures in religious contexts than scientific ones if they believe that a variety of answers may be more acceptable in religion than in science (see e.g., Liquin, Metz, & Lombrozo, 2020). A related question is whether the propensity for speculation and the acceptability of entertaining conjectures differs by domain or situational and social demands. For example, people might differ in their willingness or ability to generate useful speculations, perhaps modulated by their own prior knowledge and an estimation of their own and others' expertise.

While little is known about the origins of speculation in early childhood, it is intuitively clear that children begin attempting to answer questions as soon as they can ask them. Children's conjectures are necessarily limited by their world knowledge and are often wrong with respect to the facts. Nonetheless, a remarkable feature of children's speculations is that their answers are "at least wrong". For instance, a three-year-old with whom the senior author is well acquainted once speculated in response to the announcement that everyone had to turn off their cell phones when the plane took off that this was because "Planes are noisy and you wouldn't be able to hear if you talked on the phone." This answer is wrong but it is "at least wrong". Consider the infinite variety of things she could have said that would have been, even if factually correct, simply irrelevant (e.g., “Planes are silver and phones are silver”; “planes are big and phones are small”; see Schulz, 2012a for discussion) This early emerging ability to map the structural form of an answer onto a question might be critical to how we can generate new hypotheses. By placing high value on getting our questions answered, and by readily entertaining new ideas before obtaining evidential support, we may be motivated to explore and inquire in ways that can generate unexpected discoveries.

In conclusion, the results point to a willingness to entertain potentially explanatory but unverified speculations beginning in early childhood. Children as young as four knowingly rejected known information in favor of the unknown, but only when the conjectures addressed an otherwise unanswered question. When evidence is unavailable, children and adults willingly drop their empirical stance to consider the potential value of new ideas, even when these proposals are advanced by uncertain speakers and in tension with prior expectations, in order to find possible solutions to otherwise unresolved problems.

Chapter 4

Not playing by the rules: Exploratory play, rational action, and efficient search

Abstract

Recent studies suggest children's exploratory play is consistent with formal accounts of rational learning. Here we focus on the tension between this view and a nearly ubiquitous feature of human play: In play, people subvert normal utility functions, incurring seemingly unnecessary costs to achieve arbitrary rewards. We show that four-and-five-year-old children not only infer playful behavior from observed violations of rational action (Experiment 1), but themselves take on unnecessary costs during both retrieval (Experiment 2) and search (Experiments 3a-b) tasks, despite acting efficiently in non-playful, instrumental contexts. We discuss the value of such apparently utility-violating behavior and why it might serve learning in the long run.

4.1 Introduction

Play is one of the most charming – and perplexing – behaviors of early childhood (see e.g., Berlyne, 1969b; Bruner et al., 1976; Chu & Schulz, 2020b; Doebel & Lillard, 2023; Lillard, 2015, 2017; Lockman & Tamis-LeMonda, 2021; Pellegrini et al., 2007; Scarlett, Naudeau, Salonius-Pasternak, & Ponte, 2005; Singer et al., 2009; Zosh et al., 2018, for discussion and reviews). Play in humans and other animals may serve many biological and social functions, including acting as an honest signal of fitness (Alessandri, 1991; L. L. Sharpe et al., 2002, but see Held & Špinka, 2011), directly promoting fitness (Byers & Walker, 1995; Pellegrini & Smith, 1998), and supporting emotion regulation and promoting social bonds (Drea et al., 1996; Galyer & Evans, 2001; Gilpin, Brown, & Pierucci, 2015; Lillard, 2017; Palagi, 2008; Panksepp, 1981). Some kinds of play may also be evolutionary spandrels (Gould & Lewontin, 1979), persisting simply because behaviors that are adaptive in some contexts tend to be repeated even in non-functional contexts. (Thus, dolphins who blow bubble nets to catch fish may continue to blow bubbles when no fish are around; McCowan et al., 2000; Pace, 2000; F. A. Sharpe & Dill, 1997). These different accounts of play are not mutually exclusive, and each likely characterizes some aspects of play behaviors across species (see Chapter 2 for review). However, our primary interest here is in cognitive accounts of play, and in particular the connection between exploratory play and learning.

A broad interest in the relationship between play and learning has influenced work in developmental psychology, education, and ethology for over a century (Groos, 1901; Gulick, 1920; Montessori, 1912/1964) and today inspires both approaches to engineering autonomous agents in robotics, machine learning, and AI (Baranes & Oudeyer, 2010; Burda et al., 2019; Forestier et al., 2017; Pathak et al., 2017; see Colas, Karch, Sigaud, & Oudeyer, 2022; Oudeyer, 2018, for discussion) and a rich tradition in developmental cognitive science (e.g., Lapidow & Walker, 2020; C. H. Legare, 2012; Perez & Feigenson, 2022; D. M. Sobel, Benton, Finiasz, Taylor, & Weisberg, 2022; Stahl & Feigenson, 2015, 2018). We are advocates and contributors to this line of work and review it in detail to follow. However, play and learning are both complex phenomena and the relationship between the two remains far from simple. Here we focus on a fundamental challenge in connecting play to formal accounts of learning. We suggest that even in the context of relatively straightforward exploratory play, children tend to violate principles of efficient planning and rational action.

To begin however, we note that there is an abundant literature suggesting that children's spontaneous exploration and exploratory play is indeed sensitive to opportunities for expected information gain. Indeed, the fact that infants look longer at events that violate their expectations has been the basis for infant looking time paradigms for many decades (see e.g., Csibra, 2016, for review) and recent work has attempted to quantify the information-theoretic factors that affect infants' visual search (e.g., Kidd et al., 2012; Kidd, Piantadosi, & Aslin, 2014; Raz & Saxe, 2020; Sim & Xu, 2019).

Moreover, infants do not simply engage in rational visual search; they also selectively explore and manipulate objects in ways that are sensitive to uncertainty, prediction error and opportunities for information gain. Thus, for instance, infants selectively explore objects that appear to violate their prior expectations and intuitive theories. Thirteen-month-olds spend more time touching and reaching into a box that generated an unexpected sample (i.e., a uniform sample of colored balls from a box containing balls of many colors) than a box that generated an expected sample (i.e., balls of many colors; Sim & Xu, 2017a). Similarly, eleven-month-olds selectively bang objects if they appear to violate solidity but selectively drop them if they appear to violate gravity selectively (Stahl & Feigenson, 2015, 2018). Critically, if infants are given information that explains away the seeming violations (e.g., the solid wall that the ball seemed to pass through is revealed to be an archway with a hole in it), they no longer engage in this kind of exploration (Perez & Feigenson, 2022).

The evidence for rational exploration is even more extensive in older children. Two to five-year-olds selectively choose, design, and communicate informative interventions that disambiguate evidence in play, exploring more when evidence violates their prior beliefs (e.g., Bonawitz et al., 2012; C. H. Legare, 2012; Schulz, Goodman, Tenenbaum, & Jenkins, 2008) and also when evidence is ambiguous or confounded (Schulz & Bonawitz, 2007; Sodian, Zaitchik, & Carey, 1991; van Schijndel et al., 2015). Moreover, children systematically explore longer in contexts where there is higher uncertainty (Lapidow, Killeen, & Walker, 2022; Siegel et al., 2021) and selectively explore in ways likely to generate informative evidence for themselves and others (Butler, 2020; Butler & Markman, 2012; Cook et al., 2011; Gweon & Schulz, 2018; Lapidow & Walker, 2020). Children also learn from the evidence they generate in play (Lapidow et al., 2022; Lapidow & Walker, 2020; Sim & Xu, 2017b; D. M. Sobel et al., 2022; Walker & Gopnik, 2014).

Such findings about exploratory play are broadly consistent with work on children's early understanding of principles of rational action. Even infants expect agents to act efficiently to achieve their goals (e.g., Bálint, Csibra, & Kovács, 2021; Csibra, Bíró, Koós, & Gergely, 2003; Gergely & Csibra, 2003; Gergely, Nádasdy, Csibra, & Bíró, 1995; Liu & Spelke, 2017; Mascaro & Csibra, 2022), and toddlers and preschoolers engage in rational planning, weighing the cost of acting against its value (Bridgers, Jara-Ettinger, & Gweon, 2020; Liu, Gonzalez, & Warneken, 2019; Sommerville et al., 2018) and assuming others will do the same (Aboody, Zhou, & Jara-Ettinger, 2021; Jara-Ettinger, Floyd, Tenenbaum, & Schulz, 2017; Jara-Ettinger et al., 2015).

However, the focus on identifying aspects of exploratory play consistent with accounts of rational exploration arguably obscures some of the richness of play behavior itself. This becomes clear in considering the gap between any laboratory experiment on play and what children might do with the same stimuli if left to their own devices. Thus, for instance, four to eight-year-olds may well shake a box of marbles longer in proportion to how difficult it is to distinguish competing hypotheses about its contents (Siegel et al., 2021). However, on their own, children might never generate that specific behavior and would instead presumably engage in all manner of behaviors the details of which would be hard to predict a priori (e.g., rolling the marbles across the table, arranging them into patterns, or placing the box on top of the marbles to make a toy car). Indeed, arguably, children's play may be characterized by nothing so much as their tendency to invent novel goals and problems for themselves (Chapter 2; Chu & Schulz, 2020b).

The idiosyncratic nature of children's play has led some researchers to emphasize not the efficiency or rationality of play but its seeming arbitrariness, leading to proposals that the randomness and variability associated with play may themselves be important to learning (Dayan & Sejnowski, 1996; Gordon, 2020; Ossmy et al., 2018, e.g.,[). However, neither random behavior, nor a mere preference for doing new things, is likely to support learning in open-ended contexts where rewards are sparse (Oudeyer et al., 2007). Moreover, children at play do not simply engage in random behaviors (Meder, Wu, Schulz, & Ruggeri, 2021); they invent novel goals and plans (*"Let's go down the slide backwards"; "Let's cross the dining room without touching the floor"*; e.g., Colliver & Fleer, 2016).

We are struck by children's ability to invent new problems, goals, and constraints for themselves, and, in particular, by children's willingness to incur seemingly unnecessary costs to achieve arbitrary rewards. One might account for some kinds of costly, self-handicapping behaviors (e.g., crows and primates deliberately balancing on unstable branches) with respect to motor learning and adaptive skill building (Petrů, Špinka, Charvátová, & Lhota, 2009; Spinka et al., 2001). However, it is less clear what to make of the full range of seemingly unnecessary costs children incur in play, especially when children have already acquired the requisite motor skills (*"Let's crawl around under the bed sheet"; "Let's stick the crayons into vent"; "Let's blow the peas off the plate"*). Moreover, the gap between the complexity of play behaviors and its payoff in generalizable skills only becomes greater as play becomes more elaborate ("Let's use black and white stones to control territories on a grid"). Given the richness, variability, and arbitrariness of the problems humans create in play, we are disposed to take the arbitrariness of much distinctively human play seriously, and to consider the value that taking on new (and prima facie unnecessary) costs might have for the flexibility and productivity of human cognition.

The first step - and the focus of the current paper - is to look at whether children's exploratory play is indeed distinguished from functional behavior by the prevalence of seeming violations of rational, efficient action. Although many games are characterized by otherwise unnecessarily costly actions (e.g., jumping over chalk lines in hopscotch; picking up objects before the next bounce of a ball in jacks) to our knowledge, no previous work has investigated this phenomenon experimentally. Here we look both at whether children use apparent violations of principles of rational action to decide when others are playing (Experiment 1 and whether children themselves adopt unnecessary costs when playing. To establish the generality of the phenomena, we look at both retrieval tasks (Experiment 2) and exploration tasks (Experiment 3).

We focus on preschoolers on pragmatic grounds: They are the youngest children we can test with the linguistic and executive function skills to follow simple task instructions. Although we rely on explicit verbal instructions (to play or to achieve a functional goal), our tasks do not require or assume that children have metacognitive awareness of their tendency to incur unnecessary costs in play (See Goodhall & Atkinson, 2019; Wing, 1995, for work on children's metacognition about other aspects of play). Much as children recognize and produce grammatical sentences without knowing how, children might adopt different costs and rewards in play and expect others to do the same without explicit awareness. Note also that the current study is designed to establish whether children selectively engage in inefficient actions during play. Because it is not yet clear to what extent the phenomenon exists, we will focus here on characterizing the behavior and leave for future work any consideration of mechanisms that might underlie any observed developmental changes across ages.

4.2 Experiment 1

Our first question is whether violations of rational action might help children recognize others' behavior as playful. Of course, play is not the only possible explanation for seeming violations of rational action. Many valuable behaviors (social norms and rituals, steps necessary for effective tool use, safety procedures, etc.) are cognitively opaque (see e.g., Kenward, Karlsson, & Persson, 2011; Keupp, Behne, & Rakoczy, 2013, 2018; Rakoczy, Warneken, & Tomasello, 2008) and researchers have suggested that children's tendency to imitate even seemingly unnecessary, inefficient actions may be critical for transmitting both instrumental skills and social conventions (Horner and Whiten 2005; Keupp et al. 2018; C. Legare and Nielsen 2015; C. Legare, Wen, Herrmann, and Whitehouse 2015; Lyons, Young, and Keil 2007; Nielsen, Cucchiaro, and Mohamedally 2012; Over and Carpenter 2013; see Hoehl et al. 2019 for review). Seemingly inefficient actions can also indicate that the actor's goal was simply to perform the movements for their own sake, such as in dance (Schachner & Carey, 2013) or were intended as communicative gestures (Royka, Chen, Aboody, Huanca, & Jara-Ettinger, 2022), instead of reaching for or manipulating other objects.

However, children may nonetheless believe that violations of efficiency are characteristic of play – even relatively straightforward, goal-directed exploratory play. Consistent with this, there is some evidence that children are especially likely to imitate seemingly unnecessary means to an end in playful contexts (Nielsen, Moore, and Mohamedally 2012; Schleihauf, Graetz, Pauen, and Hoehl 2018; see Hoehl et al. 2019 for discussion). Here we ask whether, in the absence of any other discriminative cues, and given otherwise neutral behaviors (collecting sticks, retrieving a box, pushing buttons), children use a violation of efficient, rational action to decide who is playing.

4.2.1 Methods

Participants

In Experiment 1, we tested 24 children (13 females; mean age: 4.95 years, SD = 1.24, range = 37 - 83 months). The hypothesis that children would identify the inefficient actor as playing was specified ahead of data collection, but not formally pre-registered. Given our directional prediction, we chose this sample size to yield 80% power to detect a moderately large effect (Cohen's g = 0.25). Four additional participants were excluded for ambiguous responses (n=2) or not being fluent in English (n=2).

All children in this and the following experiments were recruited and tested between June 2019 and February 2020 from an urban children's museum in the United States. Each child participated in exactly one experiment. Parents provided informed consent, and children received stickers for their participation. Although we did not collect participants' demographic information, participants reflected a range of ethnicities and socioeconomic backgrounds from the museum visitors: 70% White, 3% Black, 9% Asian, 7% other races, 11% two or more races, 10% Hispanic or Latino (any race)

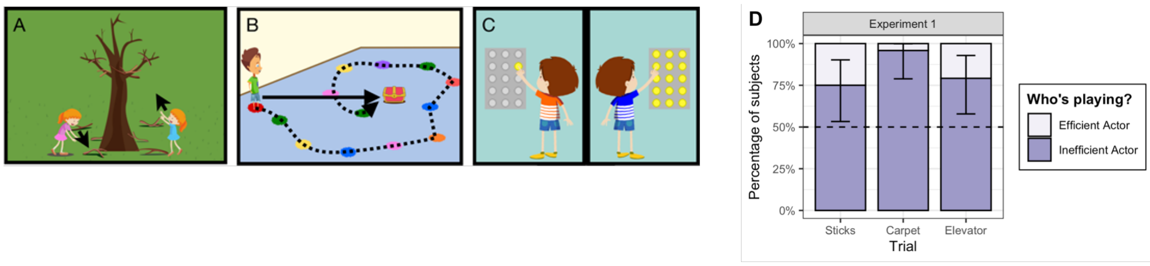

Figure 4-1: Materials and results of Experiment 1. Each story showed two characters either efficiently or inefficiently completing a task. (A) Asked to retrieve sticks to make a fire, one character reaches efficiently for sticks on the ground and the other reaches inefficiently for sticks on a tree. (B) Asked to retrieve a key from the red box, one character runs efficiently in a straight line and the other hops inefficiently in a spiral. (C) When taking the elevator to go home, one character presses just one button and the other presses all the buttons. (D) Children preferentially identified the inefficient actor as the one who was playing. Error bars show 95% bootstrapped confidence interval on the average response.

and with about 30% of attendees visiting on days when there is free or discounted admission (Boston Children's Museum, January 20, 2023).

Materials and Procedure

Participants were tested individually in a quiet room. The experimenter explained that they would watch three stories, each involving two child characters attempting the same task. She explained that "one of the children will just do what they're supposed to do" and the other child "will play". Participants were told to guess who was playing.

The stories were presented on a laptop computer in a fixed order (sticks, key, then elevator; see Figure 4-1A-C). On each trial, one character acted efficiently and the other acted inefficiently towards the goal. Each character's behavior was shown sequentially, one on the left and one on the right side of the screen, with order and side randomized at each trial. Trials 1 and 2 were animated: On Trial 1, characters retrieved sticks for a campfire. The efficient actor bent down to get sticks easily accessible on the ground; the inefficient actor jumped and tried to get sticks that were out of reach on a tree. On Trial 2, characters retrieved a key from a box in the center of a room. The efficient actor ran straight to the box; the inefficient actor hopped in a circle. On Trial 3, characters had to ride an elevator to the 14th floor. The efficient character's panel had one lit button and the inefficient character's panel had fifteen lit buttons. This trial used only a static image, removing direct cues to action efficiency but previous work suggests that preschooler can infer actions on objects from static images of their end state (Jacobs, Lopez-Brau, & Jara-Ettinger, 2021; Pelz, Schulz, & Jara-Ettinger, 2020; Pesowski, Quy, Lee, & Schachner, 2020). On each trial, after the child saw both events, the experimenter asked, "*Who was playing*?". Results (whether children selected the efficient (0) or inefficient (1) action) were coded live by the experimenter.

4.2.2 Results

Figure 4-1D shows the proportion of children choosing the inefficient actor on each trial. More than half the children chose the inefficient actor at ceiling (14/24; 58%) and all but one of the remaining children selected the inefficient actor on two out of three trials (n=9 or 38%). We predicted trial-level responses using a mixed effects logistic regression model with random by-subject intercepts. The null model (syntax: choice [1|subject]) indicates that children chose the inefficient actor more often than chance of 50% ($\beta_{intercept}$ = 1.84, OR = 6.27, 95% CI:2.35-16.7, $p < .001$). Adding a fixed effect of trial did not significantly improve the null model ($\chi^2(2) = 5.5$, p = 0.06). Adding a fixed effect of age (in months) did improve model fit ($\chi^2(1) = 4.74$, p = .03; β_{age} = 0.06, OR = 1.06, 95% CI:1.00-1.12), however, this age trend did not hold after removing the single outlier participant who always chose the efficient actor.

4.2.3 Discussion

The results suggest that, in these contexts, violating principles of rational action contributed to children's tendency to attribute a behavior as play. While some of the target behaviors - jumping in the air, running in circles and pushing many buttons - may be actions that are familiar to children as play, it is also (and presumably more commonly) the case that picking up sticks from the ground, running towards a goal, and pushing buttons are also familiar actions in play. Thus, children could not have identified playful behaviors using only surface features of the actions.

As noted, there is no one-to-one mapping between inefficient actions and play. Given a different forced-choice context, children might have been equally likely to identify the inefficient actor as "naughty" or "silly" or in pursuit of an opaque, non-obvious goal. The key point for the current purposes is simply that in the absence of any other cues, children relied on the distinction between efficient and inefficient actions to distinguish play and non-play behaviors.

However, given the forced choice context in Experiment 1, we cannot be sure whether children identified the inefficient actions as play, identified the efficient actions as not playing, or both. Additionally, the inefficient actions were arguably more salient than the efficient ones (i.e., involving longer trajectories or more activated buttons); children might have chosen the targets because they were visually more interesting rather than because they violated efficiency per se. Thus, Experiment 1 provides only suggestive evidence that children understand play as the willingness to violate efficiency and adopt unnecessary costs. Note also that this was a preliminary investigation, and the experimenter was present and not blind to condition. Although she tried to maintain a neutral expression and gaze throughout, future work should replicate the design with appropriate blinding. For the current purposes however, Experiment 1 was intended primarily as a proof of concept, informative mainly in the context of the subsequent experiments. A stronger test of the hypothesis that play is characterized by manipulated utility functions would be whether children themselves engage in unnecessarily costly actions in play. We turn to this question in the experiments that follow.

4.3 Experiment 2

In Experiment 2, we used a retrieval task (analogous to the picking up sticks and retrieving the key stories from Experiment 1) to look at whether children themselves would respect principles of rational action in non-playful contexts but violate them in play. Children were placed in identical environments with identical targets and given either functional or play instructions ("*Could you help me? Maybe you could try to get [the target]*"; "Could you play over there? Maybe you could play a game to get [the target]"). We used a within-subjects design to compare children's choices in both contexts. We predicted that children would retrieve the targets efficiently in the functional context but perform unnecessarily costly actions in the play context. This experiment was pre-registered on the Open Science Framework (link hidden for anonymous review).

Critically, however, even in the playful contexts we did not expect children's play to be characterized by random or haphazard actions. Rather, we believed that children would act efficiently with respect to their self-imposed costly actions. That is, we expected that children would perform "conditionally efficient" actions; actions that were efficient with respect to the playful goal of adopting a novel, manipulated utility function. To ensure that we could code the efficiency of children's actions reliably, we intentionally included costly affordances in the environment that children could readily exploit.

We did not include three-year-olds in this study. The affordances we included (the pencils on the wall and the dots around the spiral) were just out of reach for four to five-year-olds. For shorter three-year-olds, exploiting these affordances might have been much more difficult or impossible; thus, if younger children played differently than older children, we would not know if it was an effect of age or the ways these affordances interacted with children's age. We return to younger preschoolers again in Experiment 3 (using a design that eliminated height as a factor).

4.3.1 Methods

Participants

We aimed to test 40 participants to yield 80% power to detect a medium effect of condition (odds ratio of 2.5). Our final analyses included 38 children (24 females; mean age = 5.02 years, SD = 0.52, range = 49-69 months). Sixteen additional children participated but were replaced: Many (N=10) were excluded on a single day due to a video camera failure; six others were excluded over the course of the experiment due to experimenter error (N=2), incomplete participation (N=3), or parent interference (N=1). After coding the videos, we excluded two additional children for incomplete video clips and the pandemic prevented us from replacing these participants. Participants were randomly assigned to one of four counterbalanced conditions orders (Play or Instrumental first, and Pencils or Stickers task first).

Materials

We aimed to test 40 participants to yield 80% power to detect a medium effect of condition (odds ratio of 2.5). Our final analyses included 38 children (24 females; mean age = 5.02 years, SD = 0.52, range = 49-69 months). Sixteen additional children participated but were replaced: Many (N=10) were excluded on a single day due to a video camera failure; six others were excluded over the course of the experiment due to experimenter error (N=2), incomplete participation (N=3), or parent interference (N=1). After coding the videos, we excluded two additional children for incomplete video clips and the pandemic prevented us from replacing these participants. Participants were randomly assigned to one of four counterbalanced

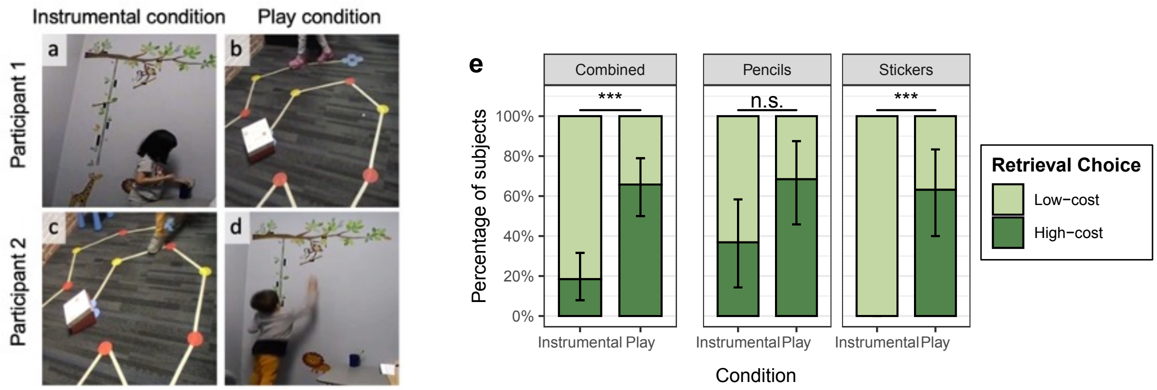

Figure 4-2: Materials and results of Experiment 2. Participants completed one Instrumental trial and one Play trial. Half the participants were assigned to the Instrumental pencils task and Play stickers task; the other half were assigned to the Instrumental stickers task and Play pencils task. Here, one child (a) efficiently retrieved pencils but later (b) took the high cost, inefficient action of walking in a spiral for the stickers. Another child (c) efficiently walked in a straight line to the stickers but then (d) took the high cost, inefficient action of jumping up to get the out of reach pencils. To best illustrate the tasks for each condition, here we presented two children who both got the Instrumental condition first and then the Play condition. However, order was counterbalanced throughout such that half the children got the instrumental condition first and half the Play condition first. (e) Our pre-registered analysis (combining tasks) found that children take low-cost, efficient actions (going straight to the target) in Instrumental conditions and high-cost, inefficient actions (taking unnecessary detours, jumping for out-of reach objects) in Play conditions. Post-hoc analyses found that in both tasks, numerically more children take high-cost actions during Play, but the effect was only significant within the Stickers task. Error bars show bootstrapped 95% confidence intervals.

conditions orders (Play or Instrumental first, and Pencils or Stickers task first).

Materials

The Pencils and Stickers tasks took place in two adjacent rooms. A third room across the hallway was used for a ten-minute distractor task (part of an unrelated experiment) in which children answered questions about short stories.

The room used for the pencil task had a colorful wall decal showing a large horizontal branch (about 6' off the ground) and a vertically hanging vine (see Figure 4-2a). Pencils were attached to this wall decal using Velcro at three different heights: 48" (taller than participants but within easy reach), 58" (requires a stretch) and 66" (requires jumping). Next to the wall decal there was a small desk (22" tall) containing a cup of pencils within easy reach of the children.

The room used for the Stickers task had a carpet on which we placed colorful dots (diameter 4") along a zigzag spiral (see Figure 4-2b). Dots were connected with straight lines of tape. The apparent start of the route was at the door where participants entered the room. A small box of stickers was placed on the dot in the very center of the spiral.

Procedure

Children completed one Play and one Instrumental trial with trial order and task counterbalanced across participants. See figure 2a-d for a schematic of the study design.

To begin a trial, the experimenter walked to the door of the appropriate room and then pretended to suddenly remember something. For Instrumental trials, the experimenter said, "*Oh, I need [some stickers /a pencil]. There's [some in that box/one over there]. Could you help me? Maybe you could try to [get that box of stickers/get a pencil].*" For Play trials, the experimenter said, "*I have to do some paperwork with your parent. Could you go play over there? Maybe you could play a game where you try to [get that box of stickers / get a pencil].*"

We took several steps to minimize the chance that adult behavior could influence the children during the tasks. The parent remained in the hallway with the door cracked open but out of sight of the child throughout both tasks. During the pencils task, the experimenter sat in the back of the room at an adult desk, with her head down, completing the paperwork until the child returned with pencils. During the stickers task, the experimenter remained outside of the room with the parent.

Our primary dependent variable was whether children performed any high-cost action as part of their first retrieval attempt. On the Pencils task, we coded low-cost behavior as retrieving a pencil from the cup and high-cost behavior as retrieving a pencil from the wall. On the Stickers task, we coded low-cost behavior as only walking or running straight to the box of stickers, and coded high-cost behavior as taking unnecessary detours (e.g. following the winding path) or performing other self-handicapping behaviors (e.g. hopping, tiptoeing, etc.). Coding decisions were made live by the experimenter and verified from video by a second coder naive to condition, with 100% inter-rater agreement on the binary judgment of high-cost/low-cost behavior. See `https://osf.io/2tde5` for video examples of the task.

Additionally, to see if children who chose the more costly actions in the Play condition acted randomly or were "conditionally efficient" with respect to the goal they chose, coders naive to hypotheses and task conditions were asked to annotate children's behavior. For children who chose to retrieve a pencil from the wall, a single coder, blind to hypotheses and conditions, judged if children deviated from a direct vertical reach by more than one hand-width. For children who approached the stickers in roundabout paths, a separate coder, also blind to hypotheses and conditions, judged whether children deviated from the spiral path the amount by which children deviated from the spiral path on the ground by more than one footstep.

4.3.2 Results

All children achieved the retrieval goal. The pre-registered analysis was the within-participant analysis, collapsing across tasks (pencils and stickers) looking at the difference between children's behavior in the Instrumental and Play conditions. As predicted, children's behavior differed in the two conditions (OR = 7.00, 95% CI: 2.09-36.65; exact McNemar's $p < .001$). The majority of participants performed low-cost actions in the Instrumental condition (n=31 of 38 children or 82%, 95% CI: 68-92%) and high-cost actions in the Play condition (n=25 or 66%, 95% CI: 50-79%; see Figure 4-2e).

We also ran three post-hoc analyses. first, we compared children's behavior across conditions within each task. In both the pencils and stickers tasks, children performed high-cost actions numerically more often when the task was presented in the Play condition than the Instrumental condition but this effect was only statistically significant in the stickers task (high-cost actions stickers task: 12/19 children (63%) in play vs. 0/19 (0%) when instrumental; fisher's exact $p < .001$; high-cost actions pencils task: 13/19 children (68%) in play vs. 7/19 (37%) instrumental, fisher's exact $p = .10$; see Figure 4-2e).

Second, we looked at whether there were any effects of order on children's performance. There were none. On Play trials, children were likely to act inefficiently whether the Play trial was presented first (n = 14/19 children or 74%) or last (n = 11/19 or 58%, fisher's exact p = .49). On Instrumental trials, children were likely to act efficiently when Instrumental trials were presented first (n=17 of 19) or last (n=14 of 19, fisher's exact p = .4). The absence of order effects was perhaps unsurprising given that the Play and Instrumental tasks took place in different rooms and the intervening ten-minute distractor task was likely to mitigate against carry over effects. However, to ensure that results were not due to children reacting primarily to the contrast in task instructions, we also ran a between-participants analysis of children's behavior looking just at performance on the first trial. Children performed significantly more high-cost actions in the Play condition than the Instrumental condition when both were the first trial (high-cost actions Play: 14/19 children (74%), high-cost actions Instrumental: 2/19 (11%); binomial $p < .001$).

finally, we asked whether children in the Play condition performed random, haphazard actions or acted efficiently with respect to the higher cost actions they'd undertaken. Of the 13 children who reached for a pencil on the wall, every child reached straight up for their target pencil; no child ever veered from this direct reach by more than one hand-width. Similarly, of the 12 children who did not run straight to the box of stickers, 11 stuck to the spiral path. Only one child ever veered more than one foot-width away from the spiral path (and she only did so for a few steps before returning to the spiral path - which she then completed for two full rounds before taking a sticker). Thus, relative to children in the Instrumental condition, children at play did incur unnecessary costs - but not by acting randomly. Instead, they appear to subvert the real-world utilities of the task environment, acting efficiently conditional on a manipulated utility function which respects additional constraints.

To enable this kind of coding, we intentionally built-in costly affordances that children could choose to exploit. However, consistent with the idea that children spontaneously manipulate utility functions in play, some children also adopted idiosyncratic costs of their own. For instance, in the Play pencils task, one child used the first pencil as a tool to swat at another pencil, and in the Play stickers task, six children went around the spiral twice before retrieving the stickers, one child went around five times (and at one point retrieved the stickers but returned them to the box before going another round!), and one child walked around the spiral backwards. In the Instrumental pencils task, 7 of the 19 children reached for the harder to get pencil on the wall, but they all immediately gave the pencil to the experimenter without taking any other actions. We did not observe any spontaneous extra costly behaviors in the Instrumental stickers task (all 19 children headed straight for the stickers).

4.3.3 Discussion

When acting instrumentally, children acted rationally and took the most efficient routes to their goals. However, even given identical extrinsic targets and environmental contexts, children adopted higher costs during play. These results support our hypothesis that children's play is characterized by manipulated utility functions in which children take actions that are unnecessarily costly given (only) the functional goals of the task. This tendency manifested across two quite different task contexts and encompassed a range of different behaviors (reaching and jumping in the pencil task, walking in a spiral in the sticker task), including some spontaneous behaviors (e.g., using one pencil to get another, repeating the spiral, walking backwards around the spiral) not particularly predicted by us as the experimenters. As expected however, when children were told to play, they did not act randomly or haphazardly. They adopted unnecessary constraints but then acted efficiently with respect to those constraints.

In this experiment, the particular actions children generated in play were (as intended) influenced by the affordances of the environment: the pencils stuck to the decal on the wall and the spiral on the ground. Superficial features of these environmental scaffolds may also account for observed task differences, such as the higher occurrence of some playful high-cost actions in the Instrumental condition of the Pencils task (compared to no high-cost actions in the Stickers task. For example, children might have found the pencils on the wall to be more novel or interesting than the patterns on a rug and thus they may have been more interested in jumping up the wall than following a spiral pattern. Alternatively, children might have found obtaining a sticker for themselves more rewarding than obtaining a pencil for someone else and thus they were more inclined to act efficiently. Critically, however, both tasks yielded a difference in performance in the Instrumental and Play conditions, and similar rates of high-cost actions on the Play trials, consistent with our proposal that adopting alternative utility functions may be characteristic of play.

4.4 Experiment 3

Experiment 2 looked at children's behavior in instrumental and playful retrieval tasks and found that children at play violated principles of rational action. However, one caveat is that children in the Instrumental condition were asked to retrieve an object that the experimenter needed (the pencil or the sticker box). Although some children may have understood that the stickers were meant for their benefit, the pencil was clearly for the experimenter's; arguably, children might act more efficiently when acting on behalf of others than when acting for themselves. In Experiment 3, we control for this possibility by asking children in the Instrumental condition to retrieve objects needed for outcomes desirable for themselves.

Additionally, Experiment 2 focused on rational action but as described in the Introduction, accounts of rational exploratory play also presume that children explore efficiently to gain information. To test the generalizability of children's tendency to take on unnecessarily costly actions during play, in Experiment 3, we asked whether children show a similar distinction between instrumental and play behavior using a search task rather than a retrieval task.

We set up tasks in which a target was equally likely to be found in either a small or large search space (i.e., among one or twelve drawers; or on a toy with two or eight buttons). If children start by searching in the smaller space, their first action will be highly informative: They would either find the target within moments of exploring, or if they fail to find the target in the small search area they will quickly know for certain that the target is in the larger search space. By contrast, if children begin by searching in the larger space, their first actions will be only minimally informative, and unlikely to be fruitful: They will have to search extensively (through many of the drawers/buttons) before either finding the target or knowing for certain that the target is in the smaller space.

Thus, if children are sensitive to the size of the search space, the rational decision is to first search in the smaller space. To our knowledge, no research has previously demonstrated that preschoolers are indeed sensitive to the size of search spaces, so this question is of interest in itself. (See Ruggeri, Sim, and Xu 2017 for work on preschoolers' sensitivity to the size of hypothesis space in question-asking tasks.) However, our primary question was not whether children would search rationally in the Instrumental condition but, presuming they do, whether they make different decisions and choose to take on unnecessary costs in play. Secondarily, as in Experiment 2, we expected that even children who violated efficiency in play would not act randomly. We thus looked at whether children who chose the high-cost tasks in the Play conditions nonetheless searched efficiently conditional on the task they had chosen.

As noted, no previous work has looked at whether preschoolers understand that it is easier to search in smaller search spaces than larger ones - let alone whether they understand that if a target is equally likely to be in each location, it is rational to begin with the smaller space. If children fail to understand this, then their tendency to violate rational action (in either condition) might simply reflect a poor understanding of the task. Although a general failure to understand the paradigm cannot, in itself, account for any condition differences that might emerge, we nonetheless thought it was important to assess whether children understood the rational decision in these tasks. Thus, at the end of the experiment, after children indicated where they wanted to start searching in the Play condition but before beginning to search, we asked children where they would search if they "*really wanted*" to find the target. If children indeed know which space is easier to search - despite selecting the more costly space in play - we can be confident that play involves taking on unnecessary costs. Note, however, that this question refers only to children's explicit understanding of the rational, instrumental action; it does not require that children have any metacognitive understanding that they violate efficient action in play. (See discussion in the Introduction.)

We ran both an original experiment (Experiment 3a) and a replication (Experiment 3b) in which we extended the study to three-year-olds. This was motivated in part by Experiment 1, in which we found no effect of age on children's judgment that play was characterized by violations of rational action, and was aided by the fact that in Experiment 3, unlike Experiment 2, the height of the child had no bearing on their ability to perform the tasks. These experiments were pre-registered on the Open Science Framework (Experiment 3a `https://osf.io/39grw`, Experiment 3b [4-5-year-olds] `https://osf.io/92dvq`, and Experiment 3b [3-year-olds] `https://osf.io/7hc9g`).

4.4.1 Methods

Participants

As in Experiment 2, we aimed to test 40 4-5-year-olds in each of Experiments 3a and 3b to yield 80% power to detect a medium effect of condition (odds ratio of 2.5) in each sample. We targeted a smaller sample of 3-year-olds because a sample of n=30 yields 80% power to both replicate the condition effect obtained in 4-5-year-olds (McNemar's OR=4.75) and to detect a large difference between the older and younger samples (binomial test, Cohen's h = 0.8).

In Experiment 3a we tested 40 children (12 female) aged four and five years (mean age: 4.87 years, SD = 0.58) randomly assigned to one of two counterbalanced task orders (Boxes or Buttons first). Two additional participants were excluded for incomplete participation (n=1) or technical failure (n=1). In Experiment 3b we tested

40 four- and five-year-olds (mean age = 4.80 years, SD = 0.53; 16 female) and 29 three-year-olds (mean age = 3.59 years, SD = 0.26; 10 female). Eleven additional participants were excluded for parental interference (n=2), exploring toys before prompts were delivered (n=2), incomplete participation (n=2), technical failure (n=1), not making a clear choice (n=1), or experimenter error (n=3). (We had aimed for 30 three-year-olds but the technical error prevented us from recovering a video file on one participant and the pandemic prevented replacing participants.)

Materials

Figure 4-3: Materials used in Experiment 3. We designed search tasks where, on each trial, a target was equally likely to be found in a smaller or larger search space. On the Boxes task, children searched for a ball in either a shelf with 12 boxes or a shelf with 1 box. On the Buttons task, children searched for a music button on a toy with 8 buttons or a toy with 2 buttons. Children completed one Instrumental trial where the target served a functional goal (to play with a ramp toy in the Boxes task or to make a robot dance in the Buttons task), then one Play trial ("hide and seek") without any functional goals. Task was counterbalanced between trials.

Figure 4-3 shows the materials used. In the Boxes task, participants had to search for a plastic ball. The larger search space was a shelf with 12 identical opaque 12"-cubic drawers (in a 4x3 array) and the smaller space was a single drawer resting atop a child-sized table. To motivate search in the Instrumental condition we used a colorful ramp toy for the ball to roll down. The experimenter and the children sat facing each other in the center of the room, equidistant from both sets of drawers. In the Buttons task, participants had to search for a button that played music. The larger search space was a round 8-button toy, and the smaller space was a 2-button toy. Both toys were diameter 12", height 3" and had colorful buttons evenly spaced around the perimeter. By design, exactly one button on each toy played music. In the Instrumental condition we used a wind-up dancing puppet to motivate search. Varying the ratio of the two options across paradigms allowed us to probe the generality of children's sensitivity to costs (i.e., whether they merely distinguished one and multiple alternatives or whether they also distinguished between fewer and more alternatives).

The Boxes task was presented identically between Experiments 3a and 3b with accompanying gestures to emphasize uncertainty about which side of the room the ball was in; children were told there was only one ball that could be in either set of drawers ("*in one drawer over there or in one of the twelve drawers over there*"). In fact, the ball was always in the drawer in the smaller search space, so that if children searched efficiently, they would find it immediately. The Buttons task was presented with slight modifications between experiments: In Experiment 3a we told children there was one music button on both toys - which was indeed the case ("*One of these two buttons makes music. One of these eight buttons makes music.*"). In Experiment 3b, to match the Boxes and Buttons tasks, we told children there was only one music button and it could be on either toy ("*Maybe one of these two buttons could play music, or maybe one of these eight buttons could play music.*"). In fact, however, we used the same toys as in Experiment 3a, so children would find the music button on whichever toy they searched.

Procedure

To control for the possibility that children might be more inclined to act efficiently when acting on behalf of others, the Instrumental tasks in Experiment 3 were for the child's own benefit: Children were asked to find a ball to slide down a ramp toy and find a button to make a robot toy dance. Given the absence of order effects in Experiment 2, the overall complexity of the design, and the fact that the overarching frame was consistent with both goals, we simplified Experiment 3 by eliminating the distractor task and having all participants complete the two trials in a fixed order: Instrumental search, then Play. On the Instrumental trial, the experimenter introduced the ramp/puppet and then announced that she wanted to play with the ramp/make the puppet dance but couldn't find the ball/didn't know which button made music. She pointed out that the target might be in either the small search space or the large search space and said: "*Can you help me find the ball/the button that plays music? It might be over there or it might be over there*". In Experiment 3a, we always introduced the side with the smaller number of options first; in Experiment 3b, we counterbalanced which side was introduced first. We coded the location of children's first search (small or larger space).

After the child found the target and used it to roll down the ramp/make the puppet dance, the experimenter transitioned to the Play trial by suggesting they play a new game with the other set of materials. The experimenter explained that she was going to hide a target for the child to find ("*We're going to play a hide-and-seek game with this ball. I'm going to hide it in a drawer, and you can look for it.*"/ "*We're going to play a hide-and-seek game with this music machine. You get to look for the button that plays music.*"; In Experiment 3b to make it clear to three-year-olds that the music button was hidden she said, "*I'm going to make one button play music and you get to find it.*") She then asked if the child would like to play in the smaller or the larger search space: "*Do you want to play with the drawers over there or the drawers over there? / With this toy or this toy?*".

After the child made a choice about where to play, but before the experimenter allowed them to go search, the experimenter asked a final question: "If you really wanted to find the ball/make music, which side/toy is easier? Looking in one drawer over there or twelve drawers over there/The two button toy or the eight button toy? Why? This allowed us to compare children's choice of where to play with their explicit judgment about which game would be easier, before children had explored themselves.

As in Experiment 2, we took several steps to minimize the chance that adult behavior could influence the children during the tasks. The experimenter and child sat equidistant between the drawer locations, and the button toys were placed equidistant from the child. The experimenter maintained a neutral expression throughout and after giving the instructions, kept her head down until the child returned with the ball/located the music button. To avoid parental interference, parents sat in a far corner of the room, reading an instruction sheet reminding them to look down and remain quiet throughout.

We coded whether children chose the low-cost action, searching first in the smaller search space (i.e., 1 drawer or 2-button toy) or the high-cost action, searching first in the larger search space (i.e., 12 drawers or 8-button toy). Coding decisions were made by the experimenter and from video by a second coder naive to condition. Inter-rater agreement was 98%; disagreements were resolved by reviewing the videos together.

Additionally, we coded whether children who chose the high-cost search task in the Play condition nevertheless searched efficiently within that larger search space. One coder blind to hypotheses and condition looked at whether children checked the twelve drawers in sequence or haphazardly; a separate coder judged whether children who played with the eight-button toy pushed the buttons around the circle sequentially or if they pressed the buttons haphazardly.

4.4.2 Results

All children searched and found the targets. The pre-registered analysis looked at the results within participants, collapsing across the Boxes and Button tasks (see Figure 4-4a). As predicted, in both the original experiment (Experiment 3a) and the replication (Experiment 3b), four- and five-year-olds chose to begin searching in the high-cost, larger search space more often on the Play trial (Exp. 3a: n=24/40 or 60%; Exp. 3b: n=30/40 or 75%) than the Instrumental trial (Exp.3a: n=9/40 (23%), OR = 4.75, 95% CI: 1.6-19.2, McNemar's p = .003; Exp.3b: n=10/40 (25%), OR = 7.67, 95% CI: 2.3-39.9, $p < .001$). Children's choices differed from chance responding: On Play trials children chose the larger search space more often than chance (Exp.3a: binomial p = .27; Exp.3b: p = .002). By contrast, on Instrumental trials, children chose the low-cost, smaller search space significantly more often than chance (Exp.3a: binomial $p < .001$; Exp.3b: p =.002). However, this is not because children failed to understand the task. Most children correctly answered that they would search the smaller space if they "really wanted to" find the target (Exp.3a: n=29/40 (73%), 95% CI: 58-85%, p = .006; Exp.3b: n=30/40 (75%), 95% CI: 60-88%, p = .002), and this was true even among children who chose to play in the high-cost, larger search space (Exp.3a: n=15/24 (63%), 95% CI: 40-81%, p = .03; Exp.3b: n=21/30 (70%), 95% CI: 51-85%, p = .04).

Among the three-year-olds in Experiment 3b, however, condition did not influence children's preference for the small or large search space (OR = 2.25; 95% CI: 0.63-9.99, McNemar's p = 0.3). Instead, the majority of three-year-olds chose the high-cost, larger search space on both the Instrumental (n=17 of 29, 59%) and Play (n=22/29, 76%) trials. Further, three-year-olds chose at chance when asked where they would search if they "really" wanted to find the target; only 12/29 (41%) correctly identified the smaller option as the easier search task.

We also did a post-hoc analysis looking at children's behavior within each task, between participants (see Figure 4-4b). In the Boxes task, four- and five-year-olds chose the low-cost smaller search space more often in the Instrumental condition than in Play (Exp. 3a: n=17/20 vs. 6/20, exact fisher's p = .001; Exp. 3b: n=17/20 vs. 5/20, $p < .001$). In the Buttons task however, this effect was weaker and only significant in the replication experiment (Exp. 3a: n=14/20 vs. 10/20, p = .3; Exp. 3b: n=13/20 vs. 5/20, p = .02). While the tasks differed in many respects, it is possible that the greater 1:12 contrast in the Boxes task provided a more compelling cost differential than the 2:8 contrast on the Buttons task.

In Experiment 3b we conducted an exploratory analysis to test the effects of age (in months) on selecting the larger search space, using a logistic mixed effects regression with condition and age as predictors and a random by-subject intercept. We found a significant age by condition interaction ($\beta = .091$, OR = 1.1, 95% CI:1.00-1.20, $p = .044$). Inspection of simple slopes within trial type found no age effect within Play trials (OR = 1.01, 95% CI: 0.95-1.08). However, there was a significant age effect in the Instrumental trials (OR = 0.93, 95% CI: 0.87-0.99), indicating that a one-month increase in age predicted a 7% lower probability of choosing the high-cost larger search space (see Figure 4-5). These results held within both the Boxes and Buttons task (see Supplemental Materials for more details).

We also looked at whether children who chose the larger search space in the Play conditions searched efficiently within this space or if they searched randomly or haphazardly. Unfortunately, only after the experiment was complete, it became clear that the camera placement obscured the full search trace on the larger set of drawers in the Play condition (although the children we could observe searched consecutive drawers). The full video data was available for the Buttons task, so we focused on this.

On the Buttons task, we identified 22 children who chose to play on the 8-button toy and who tried at least two buttons (i.e. they did not succeed on the first press). This included eight 3-year-olds and fourteen 4-5-year-olds across Experiments 3a-b. All of the eight 3-year-olds searched the buttons efficiently in order, as did most of the four and five-year-olds. Only four children pushed the buttons in a non-adjacent order: pushing two buttons simultaneously (n=1), repeatedly pressing inert buttons (n=1) or searching in arbitrary sequences (n=2). Thus, even when children chose the more costly, less efficient search space in play, they searched efficiently conditional on that choice.

4.4.3 Discussion

As in Experiment 2, the results of Experiments 3a and 3b suggest that four- and five-year-olds acted efficiently when exploring for instrumental ends but preferentially incurred unnecessary costs when playing. Note that if the larger search space was simply more salient or exciting and this affected children's exploratory behavior, four and five-year-olds would have chosen the larger search space in both conditions. Similarly, if children always preferred the fastest route to finding the target and an "easy win", they would have chosen the smaller search space in both conditions. Instead, the results suggest that four and five-year-old children choose to minimize costs during instrumental search, but selectively choose high-cost actions in play.

Critically however, and as in Experiment 2, even children who chose the high-cost search space went on to search efficiently conditional on their chosen goal: to find the target in that space. Children almost uniformly searched in an organized manner, trying each button once in order. Thus, even when voluntarily incurring unnecessary costs in play, children acted efficiently with respect to the playful utility function.

We also found a developmental effect. Three-year-olds, like older preschoolers, preferred more challenging search problems during play; however, they failed to make efficient search decisions even given instrumental goals. Three-year-olds also chose at chance (M=41%) in identifying the smaller search space when explicitly asked where they would search if they "really wanted" to find the target. As noted in the Introduction, the current study was designed to look at whether children deliberately took on unnecessary costs during exploratory play; since the phenomena had not been established, the study was not designed to look at factors that might affect changes over development. Nonetheless, we can speculate on several, not mutually exclusive, explanations for the age effect.

First, three-year-olds might have failed to represent the size of the search space as a relevant variable at all; they may have chosen the larger search space simply because it was more exciting and failed to consider the implications for efficient action. Alternatively, three-year-olds might have recognized the relevance of the size of the search space for decision-making but been unable to compare and represent the distinction in costs between the two alternatives. Changes in children's numerical cognition affect their behavior in other tasks involving costs and rewards (e.g., their preference for equal versus merit based sharing; Jara-Ettinger, Gibson, Kidd, and Piantadosi 2016); developments in children's ability to use numbers as a basis for comparison might have similarly affected children's performance in this task. In this vein, it is interesting that three-year-olds' performance was somewhat better in the task with the stronger numerical contrast (the 1 vs. 12 Boxes task vs. the 2 vs. 8

Buttons task).

Alternatively, children might have represented both the relevance of the size of the search space and the distinction in costs but failed to use this distinction as a basis for proactive planning. They may have failed to integrate all the relevant information in time to make a decision - or the larger spaces may simply have been so enticing that three-year-olds could not inhibit the desire to search there.

Note however, that when searching within a given space, three-year-olds, like older children, acted efficiently, checking buttons in order and rarely repeating actions. Thus, it is not the case that three-year-olds fail to understand efficient search altogether. This leaves open another intriguing possibility. It is possible that three-year-olds value efficiency less - or value play more - than the older participants. That is, even when given a putatively functional goal, three-year-olds might always be more likely than older children to choose to play.

Finally, as discussed, we believe these experiments are the first to show that preschoolers are sensitive to the size of a search space in goal-directed exploration and rationally prefer to search first in smaller spaces. This behavior generalized across two quite different contexts (searching for objects in chests of drawers and functional buttons on toys) and contrast ratios (1 vs. 12 and 2 vs. 8). The results suggest four and five-year-olds can anticipate and compare the relative costs of exploration and proactively select easier search problems.

4.5 General Discussion

Collectively, these results suggest that children's exploratory play is characterized by apparent violations of principles of rational action. Across three studies, preschoolers used violations of rational action to decide when others were playing (Experiment 1), and to play themselves, voluntarily incurring unnecessary costs in both retrieval (Experiment 2) and search (Experiment 3) tasks. The tendency to incur unnecessary costs in play was not due to children's failure to understand the possibility of more efficient actions: Four and five-year-olds acted efficiently in instrumental tasks even though the environments were matched across conditions. Moreover, once children decided to take unnecessarily costly actions in play (e.g., jumping for an out-of-reach target or searching in larger search spaces), children as young as three behaved efficiently with respect to these choices (jumping straight up; searching adjacent spaces sequentially). We suggest these results are consistent with the proposal that children's play is not only boundedly rational (limited by information processing constraints; Simon 1955) and resource rational (rational with respect to estimates of those processing constraints; Bhui, Lai, and Gershman 2021; Lieder and Griffiths 2020; W. J. Ma and Woodford 2020) but also conditionally rational: rational with respect to the child's self-generated utility functions.

Of course, humans are not the only species that engage in unnecessarily costly behaviors during play (Petrů et al., 2009; Spinka et al., 2001). However, the kinds of constraints and variability that non-human animals incorporate in their play are restricted and closely related to the species' behavioral and locomotor niche (e.g. somersault play is observed in Patas monkeys, which spend their time mostly in trees, but not Diana monkeys, which spend time running on the ground; Petrů et al. 2009). What's special about human cognition may not be the mere possibility of manipulating normal utilities but the flexibility with which we can do so.

In these experiments, we deliberately introduced environments that invited specific forms of play (e.g., a spiral path; pencils stuck to a tree on a wall). This made it easy not only to distinguish functional and playful behavior but also to distinguish efficient and random behavior during play (e.g., sticking close to the spiral path or vertical line of the wall versus more haphazard behavior). However, although the environmental affordances offered ready-to-hand constraints, they alone cannot account for children's behavior since children ignored these cues in functional contexts (e.g., walking straight to the stickers; retrieving the pencil from the cup). The general idea that children's behavior, even in play, is rational conditional on the self-imposed constraints they have established is consistent with other work on the ways that children respect constraints even within imaginary contexts. Thus for instance, preschoolers mop up the pretend pig who got muddy, not the one who stayed clean, and the pretend tea precisely where it spilled, not anywhere else (Harris and Kavanaugh 1993; see also Gendler 2000; Lewis 1978; Weisberg and Bloom 2009; see Harris 2021 for review and discussion).

Note however, that we do not have direct access to children's subjective utility function, and children may have imagined additional constraints beyond those that were clear given the affordances we provided (e.g. perhaps children decided to search especially slowly or quickly, or push the buttons harder than necessary). And it is of course possible that in addition to searching the larger search space in play, children could have generated additional constraints (e.g., open every other drawer first) that would have looked inefficient to us (rather than, as we observed here, conditionally efficient given the self-imposed goal). As it happens, we did not see observable cues to additional constraints here; however, future work could further investigate the nature of children's self-generated playful utility functions by asking children to provide verbal reports of their plans or by measuring play in environments with more fine-grained parameters that children may selectively manipulate.

In these experiments, we used explicit language about play throughout. In principle therefore, our results might bear more on children's understanding of the meanings of the word play, or game, than on play itself. Games - a canonical form of play - are indeed characterized by manipulated utility functions: arbitrary rewards achieved at unnecessary costs. So, perhaps children's imposition of constraints was due specifically to their understanding of what it meant to "*play a game*". However, the constraints in games are pre-established and conventionalized; whereas here children spontaneously adopted ad hoc constraints for each specific task. And although possible, we don't think there are strong grounds for believing children's behavior would differ between the instructions "*Can you play a game in here to get the stickers/pencil" versus "Can you play in here and get the stickers*" (and similarly, for "*which toy do you want to play a hide and seek game with?*" versus "*which toy do you want to play with?*") By contrast, if we had explicitly told children, "*Do anything you want in here*", we suspect that children might have played in many different ways beyond what we observed here (e.g., sticking stickers throughout the room, drawing monkeys on the tree, etc.). Such behavior would be richer and more idiosyncratic - and correspondingly harder to code - but would also involve setting costly goals. That is, we believe the particular behaviors we observed were due to the opportunities the environment afforded for play rather than the instructions to "*go play*" per se.

In this study, we constrained the extrinsic rewards by assigning target goals - and these goals were so easy to achieve that there were no obvious ways that children could relax the constraints to make the tasks any easier. Perhaps unsurprisingly therefore, in our tasks, children modified their behaviors exclusively by taking on unnecessarily costly actions. Arguably then, children's play may be characterized specifically by making an easy task more challenging rather than the ability to manipulate their utility function broadly. Clearly however, children do not always take more pleasure in more challenging tasks. (Few children would opt to walk twice as far to the school bus or search twice as many rooms for their shoes just for the pleasure of it.)

Alternatively then, children's pleasure may stem not from merely imposing higher costs but from their ability to choose the costs themselves. Indeed, when children are asked to explain what makes an activity playful, they often point to the importance of autonomy and choice (Goodhall & Atkinson, 2019). Even in playing games, which specify not only the goals but also the costs and constraints to be followed, children may take pleasure in the ability to choose and plan their actions, rather than following standard behavioral scripts or acting in the most obvious ways possible. Future work may investigate how the degrees of freedom for acting and planning shape children's decisions about which games to play and their sense of fun.

We suspect however, that we might also have induced a sense of play by allowing children to work for arbitrary rewards. We might, for instance, have compared children's responses to functional instructions: "If you go down that hallway, you can get the part we need for this balloop toy to work" versus playful ones like, "If you go down that hallway, you can get 30 balloop points". We predict that children would run faster for the arbitrary reward of "balloop points' than the functional end. Similarly, if we told children "You can put the pencils in this box here" (a functional end) or "you can put the pencils in this box with a hole in the bottom over here" (a playful end) we predict that children would opt to "clean up" the pencils in the bottomless box. Such thought experiments suggest that it's not just the willingness to incur unnecessary costs but the ability to manipulate utilities- costs or rewards or both - that children find pleasurable. Future research might look at the extent to which children in play vary their utilities broadly (e.g., by relaxing instead of imposing constraints, or by varying the reward function).

As noted, the value of setting your own goals has become increasingly clear in the fields of AI, machine learning, and robotics (e.g. Chitnis et al., 2021; Florensa, Held, Geng, & Abbeel, 2018; Gottlieb et al., 2013; Haber, Mrowca, Wang, Li, & Yamins, 2018; Kaelbling, 1993; Lynch et al., 2020; Sukhbaatar et al., 2017). Unlike agents hardwired or trained to perform pre-specified tasks, agents who set their own goals in pursuit of intrinsic rewards can learn flexibly even when extrinsic rewards are sparse (see e.g., Colas et al., 2022; Linke, Ady, White, Degris, & White, 2020; Oudeyer et al., 2007, for reviews and discussion)

However, the current results suggest that neither extrinsic nor intrinsic rewards as traditionally conceived (e.g., rewards tied to learning) adequately account for distinctively human play, even in very simple contexts like those in the current study.

We suggest that the distinctively human ability to manipulate our own utilities allows us unusual flexibility in setting new goals and creating new problems for ourselves. Rewards - even in the form of information gain - are often sparse. To the degree that we can generate reward for ourselves by creating new problems and planning and thinking within those constraints, we may be able to think of plans and ideas we wouldn't have otherwise.

We believe these results are consistent with the idea that humans not only have a remarkably flexible ability to reshape our utility functions but also find it intrinsically rewarding to do so. When children explored for instrumental ends, their utilities were determined by the most efficient way to achieve the target goals. But when children were told to play, they seemed to interpret this as an invitation to manipulate the normal utility function. Indeed, given that this difference in utilities was all that distinguished the tasks, the mere ability to manipulate typical utility functions apparently sufficed to make the task count as play.

Thus, the reward value associated with play might not be tied to learning per se but to thinking. Inventing problems we don't (actually) have might be a way of generating solutions we don't (currently) have. On this account, play is not (only) a means of gaining information. Liberated from any practical goals - even the goal of learning - play may be a means of increasing innovation. Our capacity to invent and solve small problems, and to find it rewarding, may help human learners solve a big problem: the problem of how to generate new ideas and plans in an infinite search space.

Still, why invent arbitrary utilities to generate novel ideas and plans, rather than simply explore in ways that are consistent with real world utilities and immediately likely to improve our policies and increase our knowledge of the world? One possibility is that play solves a meta exploration/exploitation problem: We can exploit our existing knowledge about valuable ways to explore (e.g., by acting efficiently to reduce uncertainty and maximize expected information gain) but we can also explore alternative ways to explore. A sure way to generate novel exploration policies is to manipulate typical utility functions by adopting unnecessary costs and trying to achieve arbitrary rewards. The world is full of unknown unknowns: if we only explored in ways consistent with expected information gain, we would miss the chance to learn the unexpected.

At this point, these ideas about the larger role of play remain speculative. But the function of play - the most characteristic behavior of our most powerful learners - has remained elusive despite decades of research. We believe there may be something to be gained by taking the seeming “uselessness” of play seriously. We hope this work contributes to asking new questions about its value.

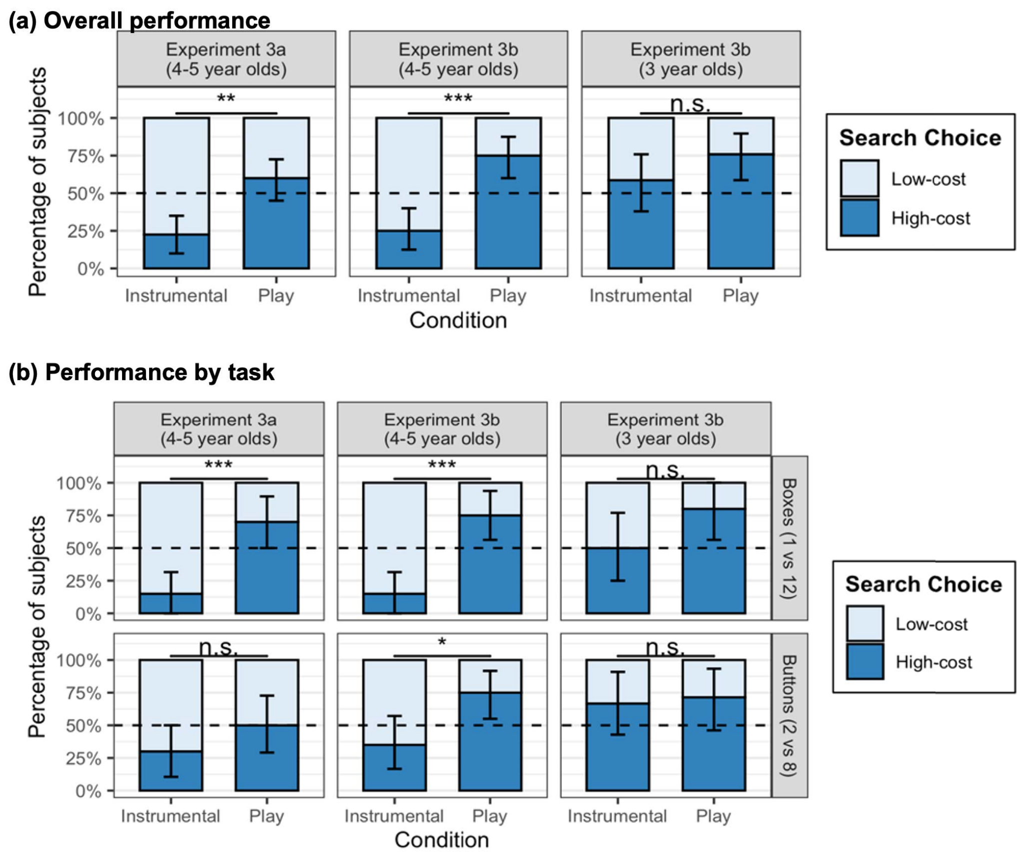

Figure 4-4: Results of Experiment 3: Children's choice of low or high-cost actions in Instrumental and Play conditions. (a) Our pre-registered analysis, finding that four and five-year-olds take low-cost, efficient actions (searching the smaller search space) in Instrumental conditions but high-cost, inefficient actions (searching the larger search space) in the Play conditions. Three-year-olds showed a similar pattern but the effect was not significant because many three-year-olds chose the larger space even in the Instrumental condition. (b) Post-hoc analyses showing that four and five-year-olds take more high-cost actions during Play than the Instrumental condition in each of the two tasks, although the effect was only significant for the Buttons task in the replication experiment. Three-year-olds a similar, non-significant trend in the Boxes task but no effect in the Buttons task. Error bars show bootstrapped 95% confidence intervals above the mean.

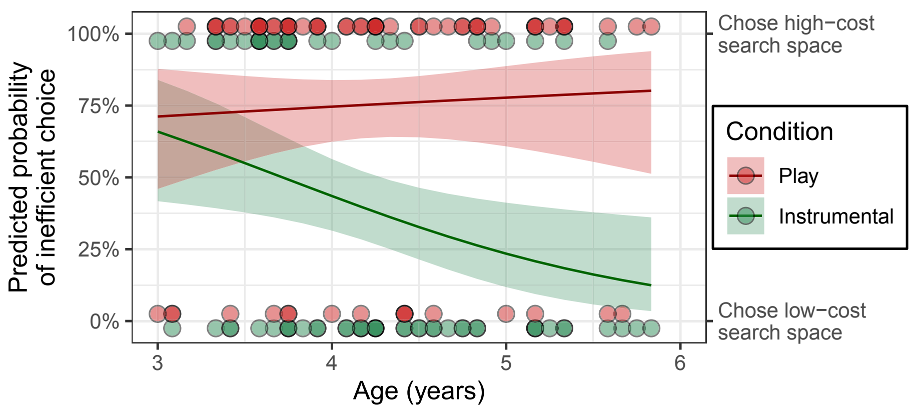

Figure 4-5: Age differences in search efficiency in Experiment 3b (N=69). Older children were increasingly likely to make low-cost, efficient choices in the Instrumental search task (green line). However, in play, children of all ages preferred the larger search space. Each circle represents the choice of one participant and lines show predicted probability of making each choice across the ages tested; shaded regions indicate 95% confidence intervals.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

"*Because I want to*": Valuing goals for their own sake.

5.1 Introduction

Both adults and young children are sensitive to the costs and rewards of actions. Cost-benefit analyses guide adults' choices not just in laboratory settings or economic decisions but also about health, crime (Becker, 1968), and sociopolitical choices about voting and alliances (Whiteley, 1995). Children also are sensitive to expected utilities: they prefer small immediate rewards to later larger ones but rationally modulate this preference according to expectations of environmental reliability (Kidd, Palmeri, & Aslin, 2013); they balance costs and rewards when exploring for information (Kidd et al., 2012, 2014; Ruggeri & Lombrozo, 2015) and they expect others to maximize utilities as well (Jara-Ettinger et al., 2015; Liu & Spelke, 2017).

This sensitivity to utilities would suggest that, given a choice between two goals of equivalent value but different costs, both adults and children should choose the goal that is easier to obtain. However, abundant research and everyday experience suggest that people do not always make the prima facie rational decision (Kahneman & Tversky, 1982). A striking instance of people's failure to maximize utilities is that after people have chosen a goal, they are often reluctant to reconsider it, even when it is clearly advantageous to reevaluate their goals and switch to an alternative (Arkes & Blumer, 1985). Investors continue pouring money into projects even once it is clear that they are unprofitable (Garland, 1990); experienced pilots continue on their flight path even when the signs of danger are evident (O'Hare & Smitheram, 1995), and doctors perseverate on treatment regimens even when better alternatives are available (Okonofua et al., 2006; L. S. Phillips et al., 2001). Indeed, faced with bad outcomes from an initial choice, people even paradoxically escalate their commitments (Brockner, 1992; Staw, 1976).

Much of the work in economics and psychology has focused on the reasons why people deviate from the predictions of rational models: misplaced optimism about the probability of success (Arkes & Hutzel, 2000); a willingness to take risks to avoid losses (Pope & Schweitzer, 2011); a sense of personal responsibility (McCarthy, Schoorman, & Cooper, 1993); social and reputational pressures (Brockner, Rubin, & Lang, 1981); or a failure to recognize alternative possibilities (Harvey & Victoravich, 2009). However, other work has focused instead on ways in which it may be rational to commit to a goal, even when seemingly better choices are available. Philosophers have suggested that tying ourselves to the mast of a thoughtful, committed decision allows us to fulfill our intentions despite temptations that might otherwise undermine our will. Among the benefits of "rational resolve" and "rational non-reconsideration" (Bratman, 1987; Holton, 2004) is avoiding the cognitive costs associated with weighing alternatives, generating new plans, and changing courses of action. Relatedly, work in psychology has suggested that apparent deviations from optimal choice can be explained by resource-rational analyses that take into account the costs of acquiring and processing information given limits on time, attention, and memory (Gershman et al., 2015; Lieder & Griffiths, 2020).

Given the myriad accounts already advanced to explain people's tendency to stick with their initial choices in the face of seemingly preferable alternatives, it might seem unnecessary to propose yet another hypothesis. However, our interest in this topic stems not from a primary interest in decision making, but from our interest in the value of goals as constraints on planning and hypothesis generation. We suggest that we may value our goals not only for their particular content or the potential reward associated with achieving them, but because goals are structured representations that support thought and action. Having a goal gives us information about which actions are worth taking and which ideas are worth thinking about. Whether those ideas and plans actually result in the achievement of the goal or not, they may be valuable: plans generated in the service of one goal can be decoupled from that aim and repurposed to other ends.

If the goals we choose gain intrinsic value as soon as we adopt them, people might tend to stick with a chosen goal at cost, and might do so even when the myriad other factors that can contribute to inertia in decision-making are unlikely to apply. This is not to say that we believe people will never change their minds: if a goal becomes meaningless (e.g., because the problem it was aiming to solve no longer exists) or if the cost differential between a chosen goal and an alternative becomes extreme, we expect people to seek out and pursue alternative goals. The idea that goals have an intrinsic value regardless of whether they are fulfilled is intended to supplement ordinary considerations of utility, not supplant them.

To empirically examine the intrinsic value we may attach to goals, we conducted a series of experiments asking whether and when people continue to pursue goals that no longer maximize utilities. In each experiment, participants completed multiple trials involving a choice between two goals with obviously different (but always achievable) action costs. Within a trial, the two goals are designed to be equally compelling, so that at baseline, each will be adopted by roughly half the participants. Experiment 1 establishes the phenomena, using a between-subjects design. In the baseline condition (Goals + Costs), participants choose between morally-laden goals ("Who do you want to help?") with full knowledge of both goals and their associated costs. In the critical test condition (Goals First), participants first choose their goal in the absence of any other information. We then reveal the costs associated with each goal, such that the chosen goal always has a higher cost, and ask if participants stick with the costly original goal or switch to the lower-cost alternative. In Experiment 2, with the same stimuli as in Experiment 1, we use a within-subjects design comparing goal persistence in the Goals First condition against a control condition (Goals Devalued). Both conditions are identical except that in the Goals Devalued condition, after choosing their goal, participants see the associated action costs but also learn that their chosen goal is no longer especially meaningful (because the problem has disappeared or because others have already solved it).We predicted that participants in the Goals First condition will stick with their original goals and complete costly actions more often than participants in the baseline (Goals + Cost) or control (Goals Devalued) conditions. finally, in Experiment 3 we replicate the between-subjects design of Experiment 1 with non-moral ("Which do you want to make?"), to minimize any extrinsic value associated with the goals.

In Experiments 4-6 we replicate and extend the initial experiments to young children (ages four to six years), using the same prompts, cover stories, and goals. We run experiments in both adults and young children for two reasons. first, although there has been a lengthy literature on stickiness in adults, we are unaware of literature on children's tendency to persist on costly goals in the face of easier alternatives. This is interesting to test in its own right because children might both be more likely to persist because they cannot represent action costs, or more willing to abandon one plan in favor of another. By presenting the same choices to children and adults, we can be confident that the relative cognitive costs for switching goals or maintaining two alternatives are negligible for adults. We focused on 4-6-year-olds because abundant evidence exists that they are sensitive to costs and rewards during decision-making (Jara-Ettinger, Gweon, et al., 2016; Jara-Ettinger et al., 2015).

Critically, our experiment is designed to mitigate against many existing explanations for why participants might stick with costly goals. This is not to minimize the importance of these factors in general, but to look at the extent to which we value chosen goals at cost even when other considerations are not at play. It cannot be the case that participants are committed to their goal because of sunk costs or loss aversion – at the moment of choice, participants have not engaged in any work towards the goal at all. Similarly, it cannot be the case that participants are uncertain about the relevant costs or unrealistically optimistic about the probability of success – the costs are transparent and although the costs are relatively higher in one case than the other, both are eminently surmountable. The participants are not subject to any group dynamics or reputational threats – the choices of goals are closely matched and arbitrary so deviating from them is unlikely to trigger threats to identity or self concept. On similar grounds, philosophical arguments about the virtues of rational resolve and resistance to temptation are unlikely to apply; both goals are virtuous and neither has any implications for the participants' well-being. Related pragmatic concerns about the costs of discarding moral goals are additionally ruled out in Experiments 3 and 6 with non-moral goals. finally, although we cannot rule out the possibility that there are always cognitive costs associated with changing plans, the task is designed to be almost trivially easy. Participants have a forced choice of two options and the difference in the costs of the two options can be seen at a glance (see Figures 5-1 and 5-3).

In such a context, we suggest that the reason participants stick with their chosen goal – despite its relatively higher cost – is that as soon as you've chosen the goal, you've reaped some of its rewards: you know what you are going to do and you know something about how you're going to do it (indeed in our simple case, you know almost everything about how to achieve it). That is, merely having the goal has set up a well defined space for thinking, planning, and acting. We suggest that in this kind of context, the default is not to engage in any reconsideration at all. Unless, as in our control condition, the goal is specifically devalued or (as in a condition whose outcome seem sufficiently certain that we need not run it) the absolute cost of achieving the initial goal makes it actively aversive, we predict that people will be inclined to ignore the cost differential and stick with harder goals.

5.2 Experiment 1

We began by comparing adults' choices on the critical test condition (Goals First), where participants were faced with action costs only after choosing a goal, against a baseline condition (Goals + Cost), where participants received action cost information when choosing their initial goal. While the same goals and costs were presented in both conditions, if merely choosing a goal makes it more likely that participants will stick to it, then participants in the Goals First condition would complete costlier actions more often than participants in the Goals + Cost condition.

5.2.1 Methods

Participants

Fifty-six adults were recruited on Amazon Mechanical Turk with the following qualification criteria: be in the United States, speak fluent English, and have a past acceptance rate of 95%. Each participant was randomly assigned to condition (29 Goals+Cost and 27 Goals First). Fifteen additional adults participated but were excluded from analysis for failing attention check questions (n=4) or self-reporting that they repeated the study or have previously seen the stimuli used (n=11). All participants were compensated $1.25 each for this 10-minute study.

Materials and Procedures

Participants completed an online survey taking approximately 10 minutes. The four test test trials began with a brief cover story describing two characters who were equally worthy of receiving help: both characters looked identically distraught, but each faced a different problem (e.g. sad kittens who were hungry vs. lost, shivering children who were cold vs wet from rain, scared monkeys trapped in fire vs. river, puppies stuck in a tree vs. on the road). Each character could be helped by performing some repetitive action; the action type was identical within a story but one task was always more effortful (e.g, clicking 5 vs 20 times, typing a short vs. long paragraph, searching a small vs. larger scene). This allowed us to compare the generality of action cost across different materials.

Figure 5-1a shows the critical differences in procedure for each condition, after the characters had been introduced. In the Goals First condition, we first presented both characters without their target actions (*"This puppy can't come down the tree to go home. This puppy can't cross the road to get home"*). Then, without presenting any action costs, we asked participants to choose a helping goal ("*Who do you want to help?*"). Next, we displayed the required actions for both goals. Critically, participants' chosen goal was always paired with the harder action and the non-chosen goal was always easier. We measured participants' choice to stick with their original goal or switch to the easier task: "*You wanted to help the puppy come down the tree, so you'll need to search this grid to find a ladder. Are you ready to help this puppy, or do you want to switch to the other option?*"). In the Goals + Costs condition, participants made a single choice of who to help after receiving full information about both characters and the required actions.

We randomized the order of stories and characters within stories, and counterbalanced which character required a harder task. In both conditions, the primary response measure was whether participants completed the easier or harder action on each trial.

5.2.2 Results

Our primary effect of interest is whether the likelihood of choosing the harder action differed by condition (Figure 5-2). To test this we conducted a mixed-effects logistic regression predicting action choice from condition, with random intercepts for subject. We obtained a significant effect of condition (likelihood ratio test $\chi^2(1)$=22.94, $p < .001$; OR=27.1, 95%CI=[5.51–134]), with participants choosing the harder task more often in the Goals First condition (M=2.89 trials of 4, SD=1.48) than in the Goals+Cost condition (M=1.14, SD=1.03).

To assess responses against chance responding, we calculated estimated marginal means per condition (i.e. model predicted probability of choosing the harder task on any given trial). Participants in the Goals First condition chose the harder task more often than chance (M=86%, 95%CI=[60–96%], $z = 2.53$, $p = .011$), with few adults always choosing the easier drawings (n=three or 11%, not different than chance of 6.25% participants). In the Goals+Cost condition however, the harder drawing was chosen less often than chance (M=18%, 95%CI=[6–43%], z = -2.38, p = .018), with 10 adults (34%) always choosing the easier drawing (significantly more often than chance, binomial $p < .001$).

5.2.3 Discussion

We found that adults preferentially persist and take on costs to achieve their initially chosen goals. Given the objective tasks demands (to help one of the characters), this additional effort was unnecessary. Indeed participants in the Goals + Costs condition preferred the easier goal. However, given participants' personally adopted goals (to help a particular character in the Goals First condition), the cost differential seemed to matter much less.

One limitation of Experiment 1 is that we cannot tell if participants in the Goals First condition might have persisted with their original goals despite the higher action cost due to some pragmatic demand from being prompted to choose whether to switch or maintain their goals. In Experiment 2 we control for this possibility by allowing participants to switch on all trials. Instead, we manipulate the value of participant's goals by either resolving the chosen goal (thus devaluing it) or leaving it unresolved.

5.3 Experiment 2

In this experiment we compare responses to the critical Goals First condition and a new Goals Devalued condition, within the same participants. The two conditions were identical to the Goals First condition in Experiment 1, with one modification: on Goals Devalued trials, after participants have chosen and goal and learned about the action costs, but before taking any action, participants learn that their chosen goal is no longer especially valuable (because the problem has disappeared). Then, participants choose whether to switch or stay with their initial goal. If participants in Experiment 1 persisted on Goals First trials for reasons other than valuing their chosen goal above and beyond its' extrinsic reward, then the same factors should apply on Goals Devalued trials and we should find no condition difference. If however participants were motivated simply by a difference in wanting after having chosen a goal, then devaluing that goal should reduce observed persistence. An a priori power analysis based on pilot data indicated that a sample of n=41 would provide 80% power to detect a medium effect size.

5.3.1 Methods

Participants

Forty-one adults were recruited via Amazon Mechanical Turk, with identical procedures as in Experiment 1. Thirteen additional adults participated but were excluded from analysis for failing attention check questions (n=2) or self-reporting that they repeated the study or have previously seen the stimuli used (n=11).

Materials and Procedures

We used the same materials and procedure as the Goals First condition in Experiment 1. However, two of the four test trials were modified ("Goals Devalued" trials) to include an additional piece of information immediately after participants chose their goal. Specifically, participants saw a captioned image describing their chosen goal already being resolved (e.g., "*You wanted to help the puppy come down the tree, so you need to find a ladder in this grid. [next page] Oh! This puppy already got help.*", with an image of the puppy coming down a ladder, Figure 5-1a). After seeing this information, participants were then prompted to choose an action to complete ("*Are you ready to help the puppy come down the tree, or do you want to switch to the other option?*").

All participants completed two Goals First trials and two Goals Devalued trials. We randomized the trial order and which two trials were selected to be presented as Goals Devalued trials.

5.3.2 Results

We conducted a mixed-effects logistic regression predicting action choice from condition (Goals First or Devalued), with random by-subject intercepts. This model yielded a significant effect of condition ($\chi^2(1) = 23.7$, $p < .001$; OR=0.11, 95%CI=[.04-.31]), with adults choosing to stick with the harder task more often on Goals First trials (M=1.38 of 2 trials, SD=0.77) than Devalued trials (M=0.73, SD=0.91).

This condition effect was robust both at the group level and within individuals. Of the 37 participants, we found that 18 (49%) participants chose to complete higher-cost actions more often on Goals First trials than Goals Devalued trials. In contrast, only three (8%) participants showed the opposite pattern.

Consistent with Experiment 1, inspection of estimated marginal means indicated that participants chose the harder task more often than chance on Goals First trials (M=79%, 95%CI=[59-90%], z=2.68, p = .007), but less often than chance in the Devalued condition (M=29%, 95%CI=[14-50%], z=-1.96, p =.0495).

5.3.3 Discussion

These results replicate and extend the findings from Experiment 1. On Goals First trials, participants stuck with their chosen goal and completed the costlier action more often than chance, just as in Experiment 1. However, on Goals Devalued trials when participants learned that their chosen goal was already resolved, participants were more likely to switch to the alternative goal. Notably, the additional information provided on Goals Devalued trials did not change either the action cost or the action outcomes (i.e. creating a large bowl of kibble still helped to make cat food for the hungry kittens, even if they were no longer hungry) but it did change the value of the problem and thus the value of the plan. This result supports the idea that the preference to not reconsider alternative goals stems more from concern with the value of one's goals than from concern about the affiliated action costs.

Integrating Experiments 1 and 2, we found that adult participants rationally consider expected utilities in deciding their goals (*i.e.*, in choosing the less costly of two goals at the baseline Goals+Cost condition) but resist switching to less costly goals once they have made a choice (Goals First condition), unless the goals are explicitly discounted due to the problem being resolved (Goals Devalued condition). We designed these experiments with emotionally charged stimuli so that both goals would be equally deserving of attention, however, the moral context of helping may have introduced additional reasons to persist. For example, participants may be managing others' impressions about themselves as being mean for deserting their initially chosen character for rather small cost differentials. While these experiments were all completed by participants on their own without any observers present, we cannot rule out that these self-concept or impression management concerns given the inherently social goals. In Experiment 3, we control for this possibility by presenting participants with non-moral goals that have no instrumental or prosocial value.

5.4 Experiment 3

In Experiments 1 and 2 we found that when participants had chosen a goal, they preferred to stick to that goal over switching to an equally valuable but less costly alternative. Here we test if this behavior generalizes to non-moral contexts, by replicating Experiment 1 with a new set of scenarios. We made two important changes in this experiment: first, we used non-prosocial goals: participants chose which of two objects to "make" (see Figure 5-1b), with the objects having no instrumental value beyond the participant's own preference. In addition, because trials were shorter, and in order to obtain more precise estimates of the effect size, we asked participants to complete 8 trials instead of four as in the previous experiments. We preregistered a target sample of n=60 based on an a priori power analysis for a medium effect of condition (condition odds ratio=2.5; `https://osf.io/9nj65`).

5.4.1 Methods

Participants

Fifty-nine adults were recruited via Prolific with the following qualification criteria: be located in the United States, speak fluent English, and have a past acceptance rate of 95%. Each participant was randomly assigned to condition (29 Goals+Cost and 30 Goals First). Two additional adults participated but were excluded from analysis for failing attention check questions (n=1) or self-reporting that they repeated the study or have previously seen the stimuli used (n=2).

Materials and Procedures

We used the same procedure as in Experiment 1, with participants completing a self-paced online survey on the Qualtrics platform. However, we updated the original test trials with new stimuli (Figure 5-1b). Instead of introducing participants to a story about two characters in need of help, we simply presented participants with two objects from the same category. For example, participants chose between kites and balloons (toys), or a lollipop and a candy cane (sweets). To make an object, participants had to perform some repetitive action; within a trial, both objects required the same type of action but one was more effortful. These tasks were similar to those in Experiment 1. Across the 8 trials, we implemented two clicking tasks, two typing tasks, two search tasks, and two sorting tasks.

5.4.2 Results and discussion

As in Experiment 1, we conducted a mixed-effects logistic regression predicting action choice from condition, with random intercepts by subject. We obtained a significant effect of condition (likelihood ratio test $\chi^2(1)$=6.93, p = 0.008; OR=4.39, 95%CI=[1.48–13.04]), with participants choosing the harder task more often in the Goals First condition (M=4.63 trials of 8, SD=3.09) than in the Goals+Cost condition (M=2.86, SD=2.13).

To assess responses against chance responding, we calculated estimated marginal means per condition (i.e. model predicted probability of choosing the harder task on any given trial). The regression model estimated that participants in the Goals First condition would prefer the harder task, although not significantly different than chance (M=65%, 95%CI=[46-80%], z = 1.59, p = .11), with two adults always choosing the easier drawings. In the Goals+Cost condition, participants preferred the easier actions (M=30%, 95%CI=[16–48%], z = -2.21, p = .027). In this condition, five adults (17%) always choose the easier drawing (significantly more often than chance, binomial $p <$ .001).

Finally, we compared responses in this experiment (non-moral choices) to responses in Experiment 1 (morally-laden choices) to test for the potential effect of task context. We did so by modeling the combined data from both experiments using mixed effects logistic regressions, and comparing 3 nested models (condition only, condition and experiment, or condition, experiment and their interaction). Using likelihood ratio tests, we found that the simple condition-only model was the best fit, with no significant additional variance explained by either the condition and experiment model ($\chi^2(1)$ = .21, p = .65) or the condition and experiment interaction model ($\chi^2(2)$ = 3.62, $p <$ .16). This result suggests there were no observed differences between both experiments in participants' choice of the higher-cost action. Instead, in both experiments we found a robust condition effect, where participants in the Goals First condition preferred to stick to their chosen goal, resulting in them completing higher-cost actions more often than participants in the Joint condition.

5.5 Experiment 4

Would the observed propensity to persist with initial goals despite higher costs generalize to younger children? On the one hand, there is abundant evidence that even infants and children are sensitive to the costs and rewards of action, both for themselves and for others. For example, children prefer small immediate rewards to later larger ones but rationally modulate this preference according to expectations of environmental reliability (Kidd et al., 2013). During exploration children will balance action and information processing demands to maximize expected information gain (Kidd et al., 2012, 2014; Ruggeri & Lombrozo, 2015; Siegel et al., 2021). Children also calibrate their efforts and degree of persistence with respect to expectations about task difficulty, persisting longer for tasks believed to be more difficult or valuable (Leonard, Garcia, & Schulz, 2020; Leonard, Lee, & Schulz, 2017; Lucca, Horton, & Sommerville, 2020). These cost-benefit analyses also guide children's choices when acting prosocially. For example, when helping others, children as young as 18 months calibrate their efforts given the task difficulty and the value of helping another (Sommerville et al., 2018), and when deciding what to teach, 5-7-year-old children will balance the ease of discovery with the costs of teaching (Bridgers et al., 2020).

On the other hand, intentional commitment to prior goals and plans requires some degree of cognitive control to maintain attention and information on the selected goal, and to inhibit any goal-irrelevant information or actions. Such executive functions typically improve over development (Diamond, 2013; Munakata, Snyder, & Chatham, 2012), with marked improvements in proactive control during the preschool years (ages 4-6) in working memory and visual attention tasks (e.g., Doebel et al., 2017). Difficulties with cognitive control could lead younger children to fail to execute on goals (e.g., Marcovitch, Boseovski, & Knapp, 2007) or to switch goals more easily (Zhai, Cheng, Moskowitz, Shen, & Gao, 2022).

In Experiments 4-6, we replicate Experiments 1-3, respectively, and test if 4-6-year-old children would universally minimize action costs or if they might ascribe additional value to goals they have chosen. Throughout the following experiments, we showed children the same prompts, cover stories, and pairs of goals as adults received. However, we adapted the study procedure and response format to be engaging and within the motor abilities of four-to-six-year-olds. Specifically, children participated in a Zoom video call with an experimenter (instead of a self-paced Qualtrics survey), and completed each trial by copying easy or harder drawings onto paper (instead of clicking or typing). Thus, for both adults and children the action costs were transparent with obvious cost differentials, but adapted for different ages and experimental platforms. Figure 5-3 shows a sketch of the stimuli and procedures for Experiments 4-6.

We pre-registered research plans for all experiments involving children. For Experiment 4, we preregistered a target sample of n=60 based on an a priori power analysis for a medium effect of condition (OR=2.5; `https://osf.io/et6gs`)

5.5.1 Methods

Participants

Sixty 4-6-year-olds (M = 5.48, range = 4.50-6.42 years) were tested over Zoom with an experimenter and given a completion certificate and $5 USD Amazon gift card for participating. An additional twenty-two children were tested but excluded for inaccurately identifying drawing difficulty during practice (n=17) or due to experimenter error (n=5). Participants were randomly assigned to conditions: Goals+Cost (n=30, M_{age} =5.48 years) or Goals First (n=30, M_{age} =5.48 years).

Materials and Procedures

Children participated in a live Zoom video call with an experimenter, lasting about 20 minutes. The experimenter displayed slides through screen share and presented children with a series of binary choices between two goals or two actions. To help children verbally indicate their choices, we always introduced the first option in a green box on the left and the second option in a purple box on the right. To minimize experimenter variability, we used pre-recorded audio clips for any prompts that children had to respond to. Otherwise, experimenters followed a standard script to transition between different trials and experimental phases. Children completed 8 trials in total: an introduction with two familiarization trials and two difficulty rating trials, followed by four test trials.

Figure 5-3 shows the overall experimental procedure. The introduction trials (Figure 5-3a) were designed to help children practice reporting their binary choices over Zoom, and to ensure the fidelity of the action cost manipulation, which involved assessing which of two drawings would be easier or harder to copy. first, as a check of general motor skill, we asked children to choose and copy one of two simple shapes (semicircle / triangle). Children made drawings on paper and held up completed drawings to the webcam. The experimenter took a screenshot of their drawing and "magically" transformed it onto the slide deck by showing the appropriate animations. The second familiarization trial involved more complex shapes (star / flower). In order to reinforce children's expectation that target drawings must be copied exactly, without embellishments or simplification, we did not transform children's first attempt and instead told them: "*Hmm, the magic didn't work this time. The magic only works if you copy exactly what you see*". The experimenter then asked the child try again, and provided either a specific suggestion for improvement (e.g. "Make sure to have exactly 5 petals") to children with inaccurate drawings, or more general suggestions to children who made more accurate drawings (e.g., "*Make it bigger*"). Children's second attempt always transformed successfully, regardless of the precise visual match.

Next, children answered two difficulty rating questions which served as inclusion criteria (Figure 5-3b). On each question, children saw two drawings differing in complexity. The first question asked children whether a semicircle or flower was harder to copy, and the second question asked whether a triangle or star was easier to copy. We excluded children who answered either of these questions incorrectly.

finally, children completed four test trials (Figure 5-3c) which used the same cover stories and pairs of goals as adults saw in Experiment 1. To achieve a goal, children had to copy a drawing of an object. We manipulated action cost by creating for each object a simple, easier drawing and complex, harder drawing. We told the children, "If you can copy a picture exactly, then, it will appear in the storybook." As with the adult participants, our primary measure is whether children chose the harder or easier action on each trial, and how this choice varied with condition.

Across participants, we counterbalanced which option was introduced first, and we presented the four cover stories in one of two possible orders. For children in the Goals + Cost condition, we additionally counterbalanced across participants which of the two goals required a harder drawing (either green-purple-purple-green or the inverse).

5.5.2 Results

As in Experiment 1, our primary question is whether participants' tendency to complete the harder action differs by condition. We conducted a mixed-effects logistic regression predicting action choice from condition, with random intercepts for subject. We obtained a significant effect of condition ($\chi^2(1)1{=}18.8$, $p < .001$; OR=10.5, 95%CI=[3.37-32.5]). Children chose the harder drawing more often in the Goals First condition (M=2.77 of 4 trials, SD=1.43) than in the Goals+Cost condition (M=1.23, SD=1.17; Figure 5-4a). Inspection of estimated marginal means indicated that in the Goals First condition, the model predicted harder drawings to be chosen more often than chance (M=77%, 95%CI = [61–88%], z = 3.02, p = .003), with only 3 children always switching to the easier drawing. In the Goals+Cost condition however, the harder drawing was chosen less often than chance (M=0.24, 95%CI = [.13–.41], z = -2.96, p = .003, with nine children (30%) always choosing the easier drawing (significantly more than expected by chance of 6.25%, $p < .001$).

We also examined potential age effects by including an additional fixed effect of age (in months). This model did not explain significant additional variance (likelihood ratio test $\chi^2(1) = 1.76$, $p = .18$). Including an age by condition interaction also did not improve model fit compared to the condition-only model ($\chi^2(2) = 3.16$, $p = .21$) or condition and age models ($\chi^2(2) = 1.4$, $p = .24$).

5.5.3 Discussion

Experiment 4 suggests that four-to-six-year-old children, like adults, tend to persist with originally chosen goals even when similarly valuable alternatives are less costly. Children did so despite recognizing and preferring the lower cost option in the Goals+Cost condition, which reflects an ability to evaluate costs and a motivation to reduce costs. However, it is possible that children might persist with goals for different reasons than adults. For instance, children might be generally "stickier" than adults, or they might not be willing to switch their minds in front of an experimenter, regardless of the value of the goal. Because the Goals + Cost condition did not present children with an option to switch, but only a single choice point, it's possible that children We test these possibilities in Experiment 5.

5.6 Experiment 5

Does children's tendency to stick with their goals instead of switching to less costly alternatives stem from placing more value on their chosen goal, or does it stem from other factors, such as not paying attention to the other goal, or a pragmatic demand? In Experiment 4, only children in the Goals First condition were prompted to choose whether to switch or maintain their goals; children in the Goals + Cost condition only made a single choice about what to do. In this modified replication of Experiment 2, we control for this possibility by allowing participants to switch on all trials. We test if children's goal persistence behavior is sensitive to the value of their goal by comparing, within-subjects, responses on Goals First trials and Goals Devalued trials. If children stick with their initially chosen goals simply due to task demands or the presence of an experimenter, they should copy the harder drawing equally often on both trial types. However, if children are able and willing to consider the alternative goal, but persist simply because they value their initially chosen goal more, then they should switch goals if their chosen goal is no longer as valuable.

In this experiment, we decrease the value of children's initially chosen goal by telling them (on Goals Devalued trials) that those characters already got help, because someone else has already completed the required drawing. We pre-registered a target sample of n=41 based on an a priori power analysis for a medium within-subjects effect of trial type (OR=2.5; `https://osf.io/5skga`).

5.6.1 Methods

Participants

Forty-one 4-6-year-olds (M = 5.51, SD = .50, range = 4.50-6.42 years) were tested on Zoom with an experimenter and given a $5USD Amazon gift card for participating. An additional twenty children were tested but excluded for inaccurately assessing drawing difficulty pictures during practice (n=18), parental interference (n=1), or experimenter error (n=1).

Materials and Procedures

We used the same child-directed materials as in Experiment 4. However, throughout the study, a confederate "Sam" was also on the Zoom call, but with their video and audio feed always off so that participants could not interact directly with them. Instead, participants saw Sam's profile picture on display, which showed a child drawing. At the beginning of the study, the experimenter introduced Sam as follows: "*Today we have another child, Sam, who will also complete this drawing game in their own room! You will see all the same stories as Sam. Hi Sam!*"

As before, children completed an introduction phase followed by four test trials. Trial one and three were presented as in the Goals First condition of Experiment 4, but trials two and four were presented as Goals Devalued trials. On Goals Devalued trials, children received an additional piece of information immediately after choosing their goal and seeing the required drawings, but before choosing which picture to copy. Specifically, they heard a phone ringing, and were told, "*Hold on, it looks like Sam also chose to help the [hungry kittens], and they already drew the [cat food]! The [hungry* *kittens] already got help.*". After this interruption (which was actually a pre-recorded audio clip), children were asked the critical prompt which was the same as in the Goals First trials: "*Are you ready to draw your [cat food] or do you want to switch to their other picture? Which do you want to draw, green or purple?*"

Across participants, we counterbalanced which option was introduced first, and we presented the four cover stories in one of two possible orders, thus counterbalancing which stories were presented as Goals Devalued trials.

5.6.2 Results

We conducted a mixed-effects logistic regression predicting action choice from condition (Goals First or Devalued), with a random by-subject intercept. This model explained significant variance beyond the null model ($\chi^2(1)$=10.9, $p < .001$), with children choosing to stick with the harder task more often on Goals First trials (M=1.39 of 2 trials, SD=0.77) than on Devalued trials (M=1.0, SD=0.92; OR=0.24, 95%CI=[.10-.60]). This condition effect was consistent when looking within participants: 13 children (32%) chose higher cost actions more often on Goals First trials than Goals Devalued trials, but only 4 children (10%) showing the reverse trend.

Consistent with the previous experiments, our model predicted that children would choose the harder task more often than chance on Goals First trials (M=81%, 95%CI=[61-92%], z = 2.85, p = .004), but no different than chance on Goals Devalued trials (M=51%, 95%CI=[30-72%], z = .06, p = .95). Exploratory analyses found no significant effects of age when added to the regression model as either a main effect or interaction with trial type (χ^2s < 1).

5.6.3 Discussion

In Experiment 4 we found that although children tended to stick to their original but costly goals on Goals First trials, they did persist less when their chosen goal was already resolved. This finding replicates the results in Experiment 2. Thus, like adults, children are sensitive not only to the cost of actions but also to the value of their goals. This also indicates that children in the age range we tested generally have the executive functioning skills to switch goals on our task – however, they don't always want to.

While our main research question centers on the difference between task conditions, participants' behavior within each condition is also noteworthy. Consider Experiments 1,2, 4 and 5, which all presented participants with helping goals. In these experiments, for both adults and young children, participants stuck to the harder task more often than chance in the Goals First condition, and chose the easier task more often than chance in the baseline Goals + Cost condition. However, on Goals Devalued trials, participants did not switch to the easier task significantly more often than chance (Experiment 2: 60% rate in adults; Experiment 5: 51% rate in children). Why not? After learning that their initially chosen goal was no longer instrumental valuable (because the problem no longer existed or because someone else had already achieved it), why did participants not switch to helping the other character, especially since it would require less effort? One possibility is that our experimental manipulation left some room for ambiguity: perhaps children did not trust that the confederate on Zoom had really helped the characters, or perhaps participants (however, given the observed condition effect, this seems unlikely). Finally, perhaps children simply value their goals more than adults. We leave these questions open for future research.

5.7 Experiment 6

In the final experiment, we return to asking if the observed goal persistence on Goals First trials generalizes to non-moral contexts even for young children. We compare children's responses in the Goals First versus Goals + Cost condition in a between subjects design. As in Experiment 4, we pre-registered a target sample of n=60 (`https://osf.io/2wdne`).

5.7.1 Methods

Participants

Fifty-seven 4-6-year-olds (M = 5.60, range = 4.42-6.50 years) were tested on Zoom with an experimenter and given a $5USD Amazon gift card for participating. An additional 16 children were tested but excluded for inaccurately identifying drawing difficulty during practice (n=12), experimenter error (n=1), parent interference (n=1), or being unable to draw the very first familiarization shape (n=1). Participants were randomly assigned to condition and ages did not differ by condition: Goals+Cost (n=28, M_{age}=5.56 years) or Goals First (n=29, M_{age}=5.57 years).

Materials and Procedures

Testing sessions were conducted via the Zoom video calling platform and lasted approximately 25 minutes. We used the same child-directed procedure and slideshow set up as in Experiment 4: participants completed an introduction phase with 4 trials that familiarized them with the study mechanics and assessed their ability to distinguish easier versus harder drawings. Then, participants completed 8 test trials.

However, we made two changes from Experiment 4. First, instead of deciding who to help, these new test trials involved non-moral goals (Fig 5-1b). On each trial, children were introduced to a scene (e.g., a park, or a beach) and chose between two objects to "make". As before, we matched the two objects to be similarly attractive. For example, they could choose between making kites or balloons to play with in the park, or between two colorful balls to make on the beach. These were the same goals that adults chose between in Experiment 3. For each object, we designed one easy and one difficult version of a drawing that children could copy.

Second, we also modified the verbal prompts to shorten the overall procedure and minimize potential pragmatic cues. In Experiment 4 and 5, upon revealing the required actions in the Goals First trials, children heard, "*You wanted to help the lost kittens, so you need to copy the rocks. We'll have somebody else copy the food for the other kittens*". The last sentence was included in case children would feel upset that the other character wasn't receiving help, and so that the interruptions by the confederate in Experiment 4 would appear more natural. However, these two factors are irrelevant to the present experiment. Thus, we removed this sentence in Experiment 6.

5.7.2 Results and discussion

We conducted a mixed-effects logistic regression predicting action choice from condition, with random by-subject intercepts. We obtained a significant effect of condition ($\chi^2(1)$= 23.7, $p < .001$; OR=13.32, 95%CI=[4.82–36.80]). Children chose the harder drawing more often in the Goals First condition (M=5.21 of 8 trials, SD=2.32) than in the Goals+Cost condition (M=1.86, SD=2.32). Inspection of estimated marginal means indicated that in the Goals First condition, the model predicted that children would choose the harder drawing more often than chance (M=71%, 95%CI = [56–83%], z = 2.61, p = .009), with no child always switching to the easier drawing. In the Goals+Cost condition however, children chose the harder drawing less often than chance (M=16%, 95%CI = [8–28%], z = -4.49, $p < .001$), with 9 children (32%) always choosing the easier drawing (significantly more than expected by chance of 0.4%, $p <$ .001).

We also examined potential age effects by including an additional fixed effect of age (in months). This model did not explain significant additional variance (likelihood ratio test $\chi^2(1) = 1.49$, $p = .22$). Including an age by condition interaction also did not improve model fit compared to the condition-only model ($p = .35$) or the condition and age model ($p = .44$).

Finally, we assessed the impact of moral versus non-moral goals on children's responses. We combined responses from Experiments 3 and 5, and as before, fitted a mixed-effects logistic regression model predicting action choice from condition and experiment with random by-subject intercepts. Model comparison using likelihood ratio test found that the condition-only model was the best fit; no additional variability was explained by more complex models containing an experiment main effect ($\chi^2(1) = 1.42$, $p = .23$) or experiment by condition interaction ($\chi^2(2) = 1.51$, $p = .47$), and importantly, the experiment by condition interaction term was not a significant predictor in the full model ($\chi^2(2) = .09$, $p = .76$; OR=1.25, 95%CI=[.29–5.23]). In summary, we found a robust condition effect (OR=11.84, 95%CI=[5.56–25.20], $p <$ .001) which did not differ between the present experiments.

5.8 General Discussion

Across six experiments, we show that both adults and young children persist with costly goals, in the face of transparently less costly alternatives, despite having sunk no action costs yet, no social pressures to maintain their choice, and arguably negligible computational costs of re-planning, given that there is only one alternative option available. Instead, the default behavior favors sticking with a chosen goal and mitigates against considering other plans, even at cost, suggesting that the goal itself is valuable. We suggest that this behavior is consistent with the idea that adults and children treat goals as valuable in and of themselves, independent of the probability or outcome of achieving it.

The current study also shows that young children, like adults, both rationally consider expected utilities in deciding their goals (i.e., in choosing the less costly of two goals in the Goals + Cost condition) and resist switching to less costly goals once they have made a choice (in the Goals First condition). Arguably though, children's reluctance to switch goals might be due to the costs associated with evaluating other options. Many studies suggest that children struggle with cognitive control and switching tasks even for simple rules that they fully understand (e.g. Traut, Chevalier, Guild, & Munakata, 2021; Zelazo, 2006). However, we suspect that task demands are unlikely to account for children's performance here. Experiment 5 suggests that many children, like adults, readily switch when their initial goals are devalued and the precipitating problem is resolved.

In Experiments 1-4, we intentionally used goals with moral and emotional content (e.g., rescuing hungry kittens or monkeys stuck in trees). We did this to try to elicit something of the authentic attachment people have to real goals in the real world. Arguably however, participants were especially loyal to these goals because they involved altruistic acts for other agents. Insofar as participants felt beholden to the particular agents they had chosen to help, they might have been particularly unwilling to consider other options. However, we found similar effects in Experiments 5 and 6 using goals with no inherent moral or affective content, suggesting that this effect is relatively robust in less social contexts. Consistent with this possibility, some recent work – in domains as neutral as navigation in 2-D grids – also finds that adults are slow to correct costly paths towards initially chosen goals (Cheng et al., 2022).

An interesting nuance to our results emerged in Experiments 2 and 5. By comparing responses to Goals First and Goals Devalued trials within-subjects, we tested the possibility that participants stuck with their original goals due to task pragmatics (e.g., a bias to repeat a previous choice when prompted), a desire to not be mean to the previously chosen character, or the possibility that children's limited executive function prevented them from switching. While we found that participants switched goals more often on Devalued trials, where their original goals were no longer needed to be resolved, than on Goals First trials, their rate of choosing the less costly goal was lower than that of participants in the baseline Goals+Cost condition. Another possibility is that participants might have committed not just to the *goal* of helping a character, but also committed to an *intended course of action* (e.g., finding a ladder in the grid, or drawing a bowl of kibble). Because we only revealed the plans *after* participants had chosen a goal, participants could not have committed to the actual course of action. Rather, they might have committed to some placeholder representation (e.g., drawing *something* that will feed the kittens). Another possibility is that participants might have engaged in spontaneous planning or mental simulations, such as imagining ways to achieve their goals. These mental computations might be treated as effort towards the chosen goal - a cognitive kind of sunk cost. Future work might disentangle precisely the computations that occur at the moment of choosing a goal, and investigate what people find rewarding or costly in that process.

Why might we ascribe intrinsic value to goals from the moment we choose them, even when the outcomes have no significant import (e.g., making an image of a lollipop vs. candy cane appear), when the actual and computational costs of switching are low, and even before we have invested any action towards the goal? One possibility is that people have an overhypothesis or general belief that goals are valuable, potentially acquired through inductive learning (N. Goodman, 1955/1983). If goals support efficient reasoning and planning compared to not having a goal, we may have experienced these cognitive benefits in previous goal-directed behaviors. We may have also learned through direct experience or observing others that persisting on a costly goal can sometimes pay off, whether in eventual goal achievement or in other side effects, such as through incidental learning that occurs even during unsuccessful attempts. Or we may have learned through social interactions that persistence and goal-commitment is generally valued in our community. These possibilities are not mutually exclusive, and future work may disentangle which of these beliefs people hold, and how these overhypotheses are acquired over development or adaptive in the kinds of decision and planning landscapes we face.

Our ideas about the intrinsic value of goals remain speculative. A task with more fine-grained, quantitative measures and graded manipulations of costs and rewards would allow us to assess the value of goals with more precision. For the moment however we will simply observe the paradox that in constraining our choices of potential actions, goals motivate us to act; thus we will side with the philosophers in arguing for the rationality of non-reconsideration.

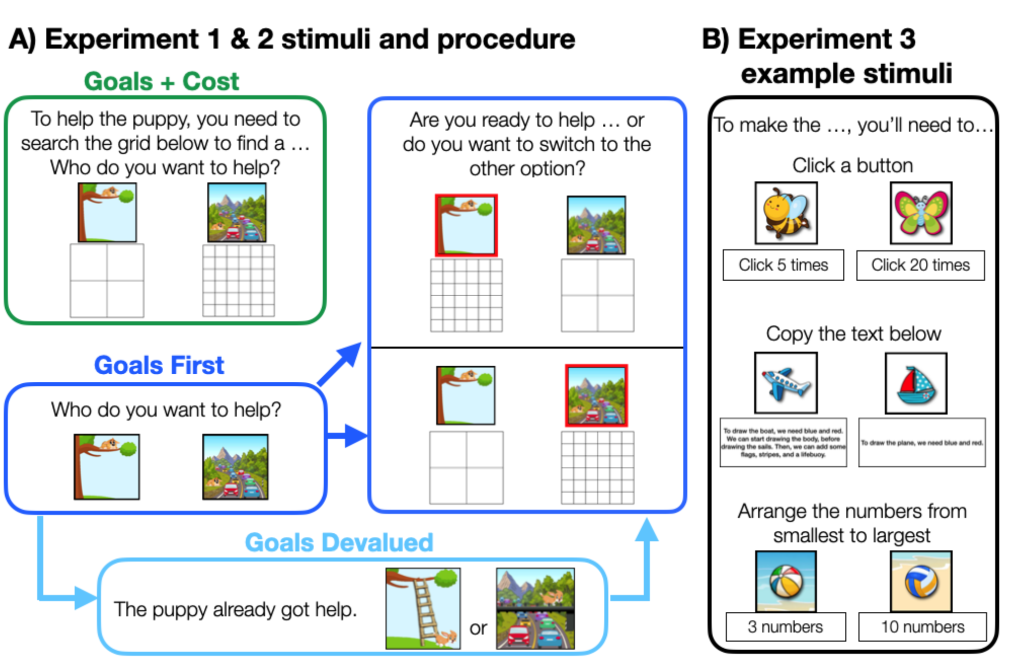

Figure 5-1: Experimental design for Experiments 1-3. **(A)** Participants in Experiments 1 and 2 completed 4 test trials, each time choosing between two goals that were matched for emotional intensity (in this example, scared puppies who were stuck on a tree vs. across a busy street). Each goal could be resolved by an accompanying action, such as searching within a grid, clicking a button repeatedly, or copying a text sample. In the baseline Goals+Cost condition (Experiment 1,3), participants saw both goals alongside their required action (one easier than the other) and made a single choice of what to do. In the critical Goals First condition (Experiments 1-3, participants first chose a goal (without knowing the action costs), before seeing that their selected goal would require a higher action cost than the alternative goal. Participants then decided whether to stay with their original goal or to switch goals (and to complete the easier action). The Goals Devalued condition (Experiment 2) was modified from the Goals First condition: after participants selected a goal, they were informed that the selected character already got help, before being asked which action they wanted to complete. **(B)** In Experiment 3, we used the same procedures as Experiment 1 but presented participants with 8 trials involving goals that did not have any moral or instrumental content.

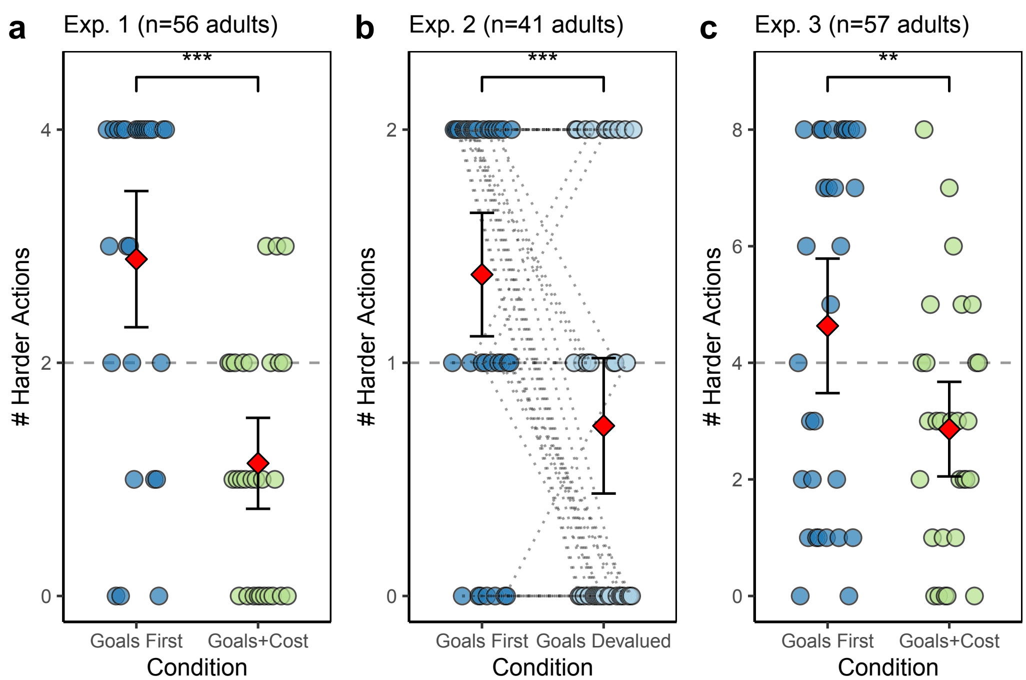

Figure 5-2: Responses by adult participants in Experiments 1-3.

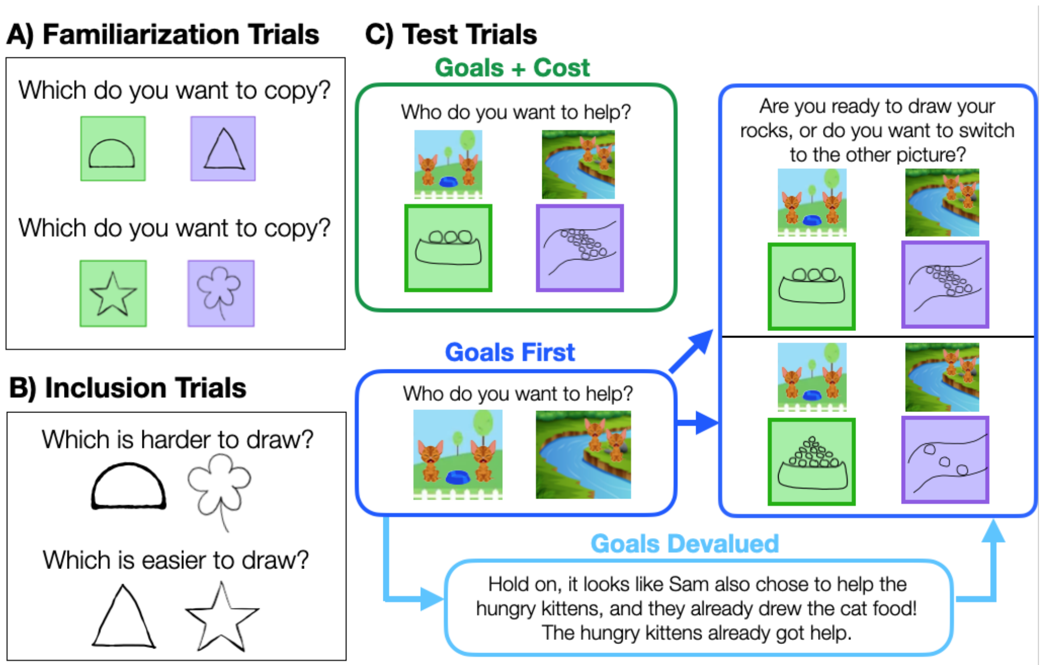

Figure 5-3: Experimental design for Experiments 4-6. **(A)** Children first completed two familiarization trials where they copied simple geometric shapes. **(B)** As a manipulation check, we asked children to complete two difficulty judgments. We excluded participants who answered either question incorrectly. (C) As in Experiments 1-3, children completed either 4 (Experiments 1-2) or 8 (Experiment 3) test trials. The overall procedure in each condition was similar for children (Experiments 4-6) and adults (Experiments 1-3), with two critical differences. first, to achieve a goal, children had to copy a drawing (as shown below each goal, in the green/purple boxes). Second, on Goals Devalued trials (Experiment 5), instead of showing an image of the completed goal, the experimenter (via a pre-recorded audio clip) explained that the other child on the Zoom call had already completed the same goal.

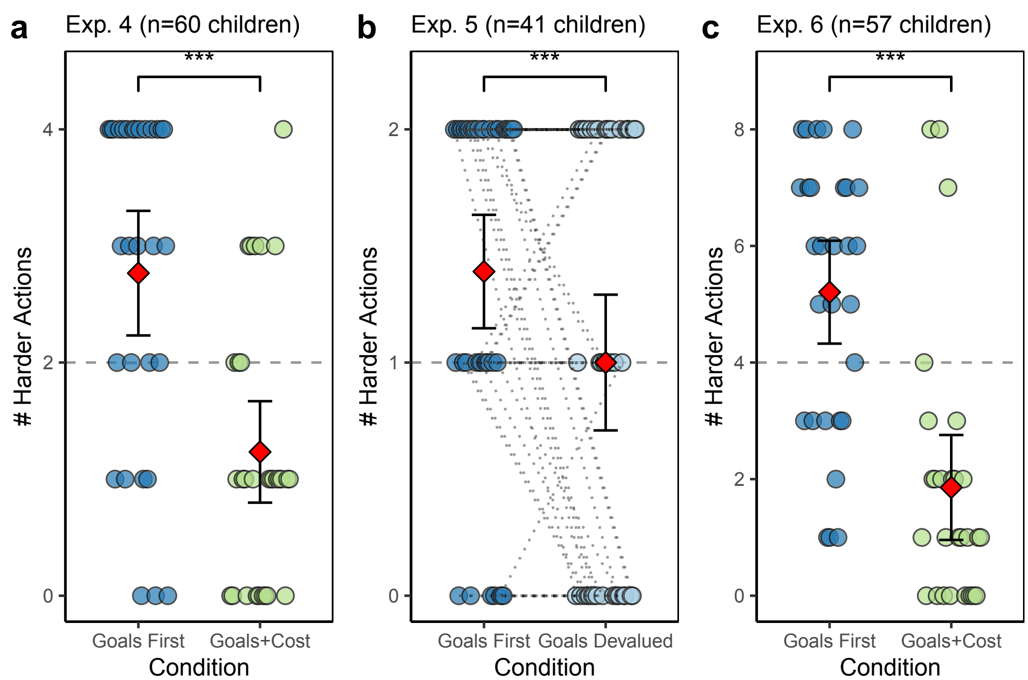

Figure 5-4: Responses by child participants in Experiments 4-6. Experiments 4 and 6 were between-subjects; in Experiment 5 we manipulated condition within-subjects

Chapter 6

In praise of folly: Flexible goals and human cognition

> It may sometimes happen that the greatest efforts of ingenuity have been exerted in trifles; yet the same principles and expedients may be applied to more valuable purposes, and the movements, which put into action machines of no use but to raise the wonder of ignorance, may be employed to drain fens, or manufacture metals, to assist the architect, or preserve the sailor.
>
> Samuel Johnson, 1751, The rambler

Abstract

Humans often pursue idiosyncratic goals that can seem remote from functional ends, including information gain. We suggest that this is valuable because goals (even prima facie foolish ones) impose valuable constraints on hypothesis generation and planning. Ideas and plans can be transmitted and adapted intergenerationally and decoupled from their original aims, leading to a proliferation of diverse ideas. If even some of these novel ideas and plans eventually pay off in learning or useful ends, this may be valuable to society as a whole.

6.1 A puzzle of goal-directed behavior

Many researchers have been interested in the extent to which humans are able to get the world right: learning so much from so little so quickly (Carey & Spelke, 1994;

Gopnik & Wellman, 2012; Tenenbaum et al., 2011). Other researchers have focused on human error, the ways that human judgments are biased and fallible (Hilbert, 2012; Schacter, 2001; Tversky & Kahneman, 1974). Recent work has bridged these accounts, considering how we might act rationally given limited cognitive resources (Gershman et al., 2015; Gigerenzer, 2008; Lieder & Griffiths, 2020). None of these approaches, however, explains humans' predilection for thinking about things that are, so to speak, neither wrong nor right: things that are imaginary, or if real, of no apparent practical value. Not only do we take on goals that have at best a tenuous connection to survival or reproductive success (e.g., digging up dinosaur bones, lining up dominoes to knock them down), we pursue goals seemingly in direct conflict with those ends (base jumping, chastity). This is of course not to say that humans are immune from evolutionary pressures. Rather, it suggests that a remarkable degree of latitude in human desires is at least compatible with, and possibly helpful to, the survival of our species. We are interested in why humans engage with goals that seem prima facie unlikely to pay off in the near term with respect to either achievement or learning. To presage our argument, we suggest that goals provide constraints on thought and support the generation of new ideas. These ideas can be decoupled from their original ends and passed on and adapted intergenerationally. If even some of these ideas do ultimately prove valuable for learning or achieving new ends, the ideas originally inspired even by seemingly foolish goals may ultimately pay off to society as whole. We begin by discussing what we mean by a goal. We then focus on four distinctive features of human goals: their flexibility, their productivity, their value to the individual, and their value to society.

6.2 What is a goal and what kinds of goals are we interested in?

6.2.1 Goals, utilities, and rational action

Evolution has equipped all organisms with ways to achieve ends useful to their survival. In the sense of the function of an adaptive behavior, we might loosely say that sunflowers have the goal of moving towards sunlight and their roots the goal of moving away from it. However, such ends are inflexible: Sunflowers cannot turn away from the sun if temperatures get too hot, and roots cannot seek daylight if temperatures get too cold. By contrast, rational, intelligent agents can evaluate different goals and act when the expected reward of achieving a target state outweighs the cost of getting there. That is, rational agents act to achieve goals with high utility. Various formal approaches in AI, machine learning, and computational cognitive science have characterized goal-directed behaviors in simple contexts (see Box 1).

Box 1: Reward-based learning and decision-making

There are many formal accounts of how agents can implement goal-directed behaviors by choosing actions that maximize future utility: the difference between the reward of achieving a goal and the cost associated with doing so. One class of approaches (Markov Decision Processes) assumes that agents make multi-step decisions within an environment containing discrete states (each associated with some reward value) as well as a set of actions (each associated with some cost) that agents can take to get from one state to another. Actions affect both the immediate reward and the probability of all future rewards, but future rewards are independent of all preceding actions conditional on the current state and action. The agent's goal is to find the combination of states and actions that maximizes cumulative total utility.

Models of reinforcement learning have been particularly influential in many recent approaches to characterizing goal-directed behavior. In model-free reinforcement learning, agents use trial and error to learn the value associated with each action and choose the course of action that maximizes total value. In model-based reinforcement learning, agents use a learned model of the environment (e.g., a spatial map, or a causal representation of the probability that actions will generate rewards) to simulate and predict possible sequences of actions and future states, summing the expected reward along these sequences.

Each approach has its advantages and disadvantages. Model-free reinforcement learning is computationally simple, but inflexible; an agent that has learned the value of actions by trial and error cannot easily learn a new policy. Model-based reinforcement learning is more flexible: An agent who has a spatial map or causal representation of actions and outcomes can simulate the consequences of changes in the model. However, it is computationally demanding to simulate the forking possibilities of state-action-value combinations in order to find high-value action sequences. Both approaches have generated valuable insights into how the brain computes reward in simple decision-making tasks (Dayan & Niv, 2008; D. Lee, Seo, & Jung, 2012). However, the assumptions underlying these algorithms may fail in the real world, where data is relatively sparse, and rewards are decoupled from near-term action consequences (Gershman & Daw, 2017). Moreover, humans seem to find value in constructing models of the world that are known to be false, and even models that will never be true. To our knowledge, no current computational approach accounts for this kind of goal-directed behavior.

6.2.2 Goals as mental states

This minimal definition of goal-directed behavior does not require that an agent represent themselves as having a goal, only that they act in ways that increase the probability of achieving it. A richer sense of goal-directed behavior is one in which goals are mental states: intentions to achieve outcomes. Goals are hierarchically structured and even animals with explicit mental state goals (for a recent account of the phylogeny of agency, see Tomasello, 2022) will have explicit access only to some levels of this hierarchy (e.g., we may be aware of the goal of picking up a pen but not the detailed subgoals in the motor plan). Here we are interested in the narrow swath of goals to which agents have conscious access; we believe these are the kinds of goals that play a critical role in thinking and planning.

6.2.3 Distinguishing goal-directed actions from other behaviors and distinctively human goals

We can twirl a pen unconsciously, as a nervous habit, or intentionally, looking for a way to open the cap; as with all mental state inferences, the behavior underdetermines the underlying mental states. Nonetheless, we can often distinguish goal-directed behavior from other behaviors empirically. Goal-directed behavior is not only efficient with respect to the goal (i.e., maximizes utilities), it is also equifinal: a goal-directed agent will find other means to pursue their end if the original path is blocked (Gergely et al., 1995; James, 1890; Uller, 2004). To borrow an example from William James: iron filings and Romeo may both move directly towards a target, but the iron filings will be stopped by a barrier; Romeo will find a way around it. Human goals lie on a continuum with that of other animals. Many animals have some flexibility with respect to their goals (see Box 2). However, while human goals have much in common with that of other species, they are also distinctive in a number of ways that we believe not only result from, but contribute to, the sophistication of our cognition broadly. We turn now to these characteristics of human goals.

Box 2: Goals in humans and other animals

At its simplest, all organisms capable of classical or operant learning can assign value to an otherwise neutral state or action if it is paired frequently or powerfully enough with a species-specific reinforcer; the animal then learns to approach or avoid the arbitrary cue. But many animals are also capable of more sophisticated forms of learning and decision-making, involving multi-step sequences (Dall, Giraldeau, Olsson, Mcnamara, & Stephens, 2005; Hunt et al., 2021). Recent models of reinforcement learning [Box 1] have characterized ways in which reward can be propagated backwards from end goals to sub-goals such that actions and states with no inherent value in themselves become valuable as steps towards desirable outcomes.

Beyond responding to reinforcers in their environment, non-human animals also act for social (Sato, Tan, Tate, & Okada, 2015) and epistemic ends (Bromberg-Martin & Hikosaka, 2009). In particular, it has been clear for decades that many animals will explore novel, complex, and surprising stimuli in the absence of any external payoff (Berlyne, 1966). Recent work has added precision to these claims, distinguishing, for instance, the impact of extrinsic reward versus reducing uncertainty on rhesus monkey's visual saccades (Daddaoua, Lopes, & Gottlieb, 2016). More broadly, computational work (Colas et al., 2022; Linke et al., 2020; Oudeyer & Kaplan, 2009) suggests that many different kinds of proxy goals support learning new knowledge and skills, including exploring rare, novel, or surprising events (Barto, 2013), trying to maximize expected information gain or the rate of reduction of prediction error (Houthooft et al., 2016; Lopes, Lang, Toussaint, & Oudeyer, 2012; Pathak et al., 2017), or trying to perform particular actions or reach particular states (Schmidhuber, 2010, 2013). Thus, even beyond the dazzling array of "hard-wired" goals innate to specific species (e.g., building hives or dams Naiman, Johnston, & Kelley, 1988; Seeley, 2014; migrating north and south Emlen, 1975; Reppert, Gegear, & Merlin, 2010), non-human animals pursue a wide range of behaviors that allow them to navigate dynamic environments and respond in real-time to change in their conditions.

All of these kinds of goal-directed actions are of course present very early in humans as well. Babies visually explore objects from birth (Amso & Johnson, 2006; Slater & Morison, 1991); manipulate objects as soon as they can reach (Rochat, 1989), and by ten months, enact simple plans to achieve rewarding goals (e.g., pulling on a blanket to get an out-of-reach toy (Willatts, 1999). Infants also actively pursue social rewards, smiling and cooing in response to loved ones (Jones, Collins, & Hong, 1991; Ruvolo, Messinger, & Movellan, 2015), responding to others' emotional expressions

(Ruba & Repacholi, 2020; Y. Wu & Gweon, 2022), and engaging in turn-taking social interactions (Murray & Trevarthen, 1986). Finally, of course, infants engage in explicitly epistemic behaviors: looking longer at rare and surprising events (Aslin, 2007; Sim & Xu, 2019), exploring objects that violate their prior beliefs (Perez & Feigenson, 2022; Stahl & Feigenson, 2015) and selectively attending to the object of adults' gaze and points (E. Y. Kim & Song, 2015). However, while humans share many goals with other species, human goal-directed behavior is also distinctive in a number of respects we discuss here.

6.3 Flexibility

Most work on distinctively human cognition has emphasized the sophistication of our representational system: its capacity for recursion (Ferrigno, Cheyette, Piantadosi, & Cantlon, 2020; Hauser, Chomsky, & Fitch, 2002), compositionality (N. A. Goodman, Tenenbaum, Feldman, & Griffiths, 2008), symbolic manipulation (Forbus, Liang, & Rabkina, 2017), communication (Premack, 2004; Tomasello & Rakoczy, 2003), etc. We have no doubt that capacities like these are fundamental to the kinds of goals humans can entertain. However, we suggest that it is not just the relative advantages of our representational system, but the relative independence of our motivational system that contributes to the successes of human cognition.

6.3.1 Long-range goals

Although many animals pursue long-term goals (migrating, building nests or dams, etc.), such behaviors are automatic and inflexible. The ability to plan towards novel future goals, even on relatively short time horizons (e.g., in selecting tools for use in a subsequent task or tokens for future bartering) has been documented only in corvids and great apes (Correia, Dickinson, & Clayton, 2007; Kabadayi & Osvath, 2017; Mulcahy & Call, 2006; Raby, Alexis, Dickinson, & Clayton, 2007, though see Redshaw, Taylor, & Suddendorf, 2017). Humans seem to be unique in our ability to set flexible goals that may take months, years, or even lifetimes to achieve. The ability to engage in long-range planning requires the ability to break a goal into sub-goals. Such sub-goals may be very far from the final goal, but the individual must still be motivated to take the first steps. This kind of planning poses a challenge to current models of reinforcement learning: Rewards depend on long-term consequences in ways that violate assumptions of the models (Gershman & Daw, 2017). However, despite uncertain and unlikely payoffs, humans routinely do engage in long-range planning, suggesting a remarkably flexible capacity for accruing intrinsic reward. We suggest that the capacity to experience reward decoupled from near-term outcomes contributed to the flexibility of human motivation.

6.3.2 Within-species variability

Individual humans differ not only in our temperament and abilities (like individuals of many species Réale, Reader, Sol, McDougall, & Dingemanse, 2007) but also in our interests: the particular things we want to do. Considerable attention has been paid to our species' ability to cooperate, divide labor, and collaborate (e.g., Tomasello & Carpenter, 2007), however, we suggest that individuals' motivation to pursue differing ends is at least as fundamental to our species' success. Specialized interests emerge early, and early interests are both variable and enduring (e.g., one study found that more than half of four to six-year-olds reported a idiosyncratic interest in a conceptual domain - ranging from bugs to ballet - and a fifth of respondents maintained that interest over two years Alexander, Johnson, Leibham, & Kelley, 2008). This variability of interests manifests in adulthood as well (Savickas & Spokane, 1999) and emerges cross-culturally (Day & Rounds, 1998; Lubinski, 2000). Thus collectively, we can pursue a remarkable range of goals; each of us is likely to invent different problems, motivated by our own particular constellation of experiences, abilities and interests.

Box 3: Many functions of play

Play is a rich, multifaceted phenomenon that likely serves different functions in different contexts. This article is on human goals broadly, and it is beyond our scope to provide a comprehensive review of the literature on play (for recent reviews, see Andersen, 2022; Chu & Schulz, 2020b; Lillard, 2015; Lockman & Tamis-LeMonda, 2021). Briefly however, researchers have suggested a variety of accounts to explain the value of play across species, including that it might serve to promote or signal physical fitness (Alessandri, 1991; Bekoff, 1988; R. M. Fagan, 1981; Held & Špinka, 2011); strengthen social bonds (Palagi, 2008; Pellis & Iwaniuk, 2000; P. K. Smith, 1982), or help animals gain information about their environment and their own competencies (Pellegrini et al., 2007; Spinka et al., 2001; Vasconcelos et al., 2015). Extensive work has also looked specifically at the potential benefits of imaginary play (Lillard, 2017; Lillard et al., 2013), suggesting that it may contribute to competencies ranging from language skills (Quinn, Donnelly, & Kidd, 2018) and counterfactual reasoning (Gopnik & Walker, 2013), to social cognition (Hughes & Dunn, 1997; Weisberg, 2015) and executive function skills (Doebel & Lillard, 2023; Lillard, 2017, but see Lillard et al., 2013).

These accounts of play are not mutually exclusive and there is good reason to believe that each characterizes some aspects of play in humans and other animals. Thus although here we argue that some forms of play, especially in older children and adults, may not be closely tied to immediate functional ends, other forms of play in humans and other species may indeed provide direct benefits to the individual.

6.3.3 Play in older children and adults: Making up problems for fun

Beyond individual differences in interests, we also invent and invest in novel goal-directed behaviors. Variability is so characteristic of children's play that some researchers have suggested that the randomness and variability associated with play might itself be important to learning (Dayan & Sejnowski, 1996; Gordon, 2020; Ossmy et al., 2018). However, neither random behavior, nor a mere preference for doing new things, is likely to support learning in open-ended contexts where rewards are sparse (Oudeyer et al., 2007). Moreover, children at play do not simply engage in random behaviors; they invent novel goals and plans ("Let's balance cups on the cat"; "Let's cross the room without touching the floor"; "Let's pretend to be rocks"; (Colliver & Fleer, 2016)). We suggest that in its variability, structure, and value, this kind of play shares much in common with aspects of human goal-directed behavior that continue throughout the lifespan.

Many animals other than humans play, and play may serve many useful ends, both cognitive and non-cognitive (see Chapter 2 and Box 3). However, starting in middle childhood, humans also sometimes engage with made-up problems and goals in ways that can seem strikingly decoupled from such outcomes.

We have in mind the idiosyncratic but perfectly ordinary activities of a six-year-old, racing through the house, who, when asked to account for his behavior, explains that he and his sloth friend are trying to put out a fire on Jupiter. This is clearly intentional, goal-directed behavior and the child may pursue it all morning without external encouragement (indeed, even in the face of some active discouragement). The child may well experience the rewards associated with simulating progress towards his imagined goal, but he is just as likely to thwart his own progress as to pursue it: He may move the goalposts on the very brink of success - making the fire leap over the hose the sloth is holding, or "realizing" that the fire is an enchanted fire, impervious to his attempts. That is, the child can vary the problem at will, and can assign costs and rewards as he likes. Moreover, the child may then abandon this goal (and any attendant mess) after a few hours and never look back. It is not that the child is not attached to his goal (woe betide the grownup who tries to interrupt a playing child) but that the accomplishment of the goal, and even measurable progress towards it, does not seem to be essential to what the child finds rewarding.

Play is not limited to children – nor is the propensity to be absorbed by arbitrary goals limited to those endeavors we call play. Much of what adults find rewarding also involves engaging with problems we don't (otherwise) have. Above we used an example from pretend play but goals can be "made-up" without being imaginary. A single salient example may suffice. A gentleman named Gareth Wild recently made headlines for successfully parking in all 211 parking spots at the Sainsbury grocery store in Bromley, England (Yuhas, 2021). The project took him six years. Although this did - remarkably - result in a moment of glory, it would hardly have been rational for him to have pursued the goal with that expectation in mind: Obsessing over parking spots in a grocery store is an unlikely route either to improved social status or to new knowledge and skills. Presumably the value of the goal is simply that Mr. Wild assigned value to it. That is, the capacity to invent and engage with arbitrary problems characterizes play in both children and adults and is sufficiently rewarding that it is what humans do for fun.

This is a compelling issue for accounts of human cognition in that we believe there are few things humans find as rewarding as setting arbitrary goals and trying to solve for them. Modern Western culture may take a particularly indulgent view of play, but play is a cross-cultural universal (Box 4). As noted, researchers have suggested many possible valuable outcomes of different forms of play in humans and other animals. However, to our mind, a striking aspect of some intrinsically motivated behavior, especially in older children and adults, is not what we could be learning or achieving, but the possibility that we may not have to experience motor learning, social bonding, progress towards a useful end, or information gain to find the activity rewarding. We will return later to the nature of what humans do find rewarding in such pursuits. Here we simply want to emphasize the paradox that humans are able to commit passionately to goal-directed activities while simultaneously remaining oddly indifferent to the functional consequences.

Box 4: Play across cultures

Play is a cross-cultural universal, even down to some of its particulars. Despite children's reputation for having short attention spans, children's play often consumes a substantial portion of their time - roughly three hours a day in the United States and Japan [134]- comparable to estimates for four to sixteen-year-olds in two subsistence communities (each quite different from the other) in the Congo basin, the Aka and Ngandu [135]. And roughly a third of playtime in both these communities is categorized as "pretense" or "idiosyncratic", comparable to estimates of pretend play in kindergartens in Eastern Slovakia [136]. Indeed, imaginative play persists even in cultures, like traditional Mennonite communities, that actively discourage it [137,138]. Researchers have documented dedicated areas for children's play (e.g. tree houses, miniature playhouses) across the world [139] and even in the prehistoric archaeological record [140], suggesting the prevalence of play throughout history.

In the sections to follow, we will suggest that our ability to flexibly decouple the rewards of goal-directed behavior from near term outcomes is beneficial because the structure of goals and problems allows us to bootstrap new plans and ideas (see Section 6.4 on Productivity). Speculatively, we suggest that the epistemic reward signals to which humans are attuned may include not only progress towards a goal, and the rate of uncertainty reduction or learning progress (Gottlieb et al., 2013) but our degree of engagement or rate of thinking (see Section on 6.5Value to the Individual). Finally, we will suggest that the richness and flexibility of human goal-directed behavior does not have to pay off in any other rewards to the individual, provided it does ground out in payoffs to society at large (see section on Value to Society).

6.4 Productivity

Flexibility, untethered from the imperative to accomplish any end, might seem inherently unproductive. Indeed, it might seem like such unproductivity is what we are arguing for. To the contrary, we believe that distinctively human goals are immensely productive - just not exclusively for extrinsic reward, social rewards, goal achievement, or even learning for the individual. We are struck by the fact that a goal, any goal - however ludicrous, unattainable, or fictitious - contains structured information that imposes valuable constraints on thinking and planning. This matters in two respects.

6.4.1 Conditional rationality

first, although people may willingly incur otherwise unnecessary costs for arbitrary and even unachievable rewards, they nonetheless behave rationally with respect to their goals. A child may fight an imaginary fire, but belying the surface irrationality (i.e., expending energy racing towards a non-existent target), the child will engage an efficient action plan, taking the shortest path to the supposed flames (narrowly avoiding any obstacles, visible only to him, along the way). In recent work, we have investigated this phenomenon experimentally (Chu & Schulz, in press). Although children violate principles of rational action in play - willfully pursuing fixed rewards at unnecessary costs - they behave rationally with respect to those goals; taking the most direct route consistent with their self-imposed constraints (see Figure 6-1a, adapted from Chapter 4). That is, consistent with abundant work on the ways children respect constraints even within imaginary contexts (e.g., mopping up the pretend tea in the precise location where it spilled; washing the pretend pig who fell in the mud, not the one who stayed in the pen; Gendler, 2000; Harris, 2021; Harris & Kavanaugh, 1993; Lewis, 1978; Weisberg & Bloom, 2009), we suggest that children's play is conditionally rational: rational with respect to the goal, problem and constraints they have set.

6.4.2 Structured problem spaces and "in-principle" solutions

Second, goals set up structured problem spaces: they provide information about how to achieve them. Suppose for instance, the child announces that the sloth is afraid to slide down the fire pole. By many measures, this is an insoluble problem by virtue of not being a problem at all: There is no sloth and there is no pole. Nonetheless, people are perfectly capable of generating candidate solutions (e.g., bribe the sloth with candy bars, tell the sloth he'll win the bravest animal of the year award, etc.). Like the goal, the solutions are at once nonsensical and reasonable: They satisfy the abstract constraints imposed by the goal (providing an unwilling agent with reasons to perform an action). By contrast, many real-world facts about sloths (that they have poor eyesight; that they can turn their heads 270 degrees) do not.

What is true for the child and sloth is, we suggest, true for the arbitrary goals that humans set in general. Both adults and children endorse speculative conjectures that provide "in principle" solutions to problems, even if they have no other evidence for them - and both adults and children prefer even highly improbable conjectures to verifiable facts if the former satisfies the constraints of the problem and the latter does not (Figure 6-1b, adapted from Chapter 3). That is, we evaluate ideas with respect to their utility - how well they satisfy the goal (i.e., the explanatory or other constraints imposed by the problem Lombrozo, 2016; Schulz, 2012a) - not just with respect to their probability. Thus, we can say of an idea "it is a good idea" not only long before we have tested it but sometimes even when we know that it is false (e.g., as when we admire the logic behind a child's speculation that boy babies grow in fathers' bellies and girl babies in mothers').

In the sloth example, we expressed both the problem and solution in language – and indeed, language offers us a great deal in the way of well-structured hypothesis spaces. For instance, long before we can answer the question, we know what kind of a response will count as answer ("who" questions refer to social networks; "where" to spatial maps; "when" questions to a timeline; "what" to a category structure; "which" to a Venn diagram; "how" to a circuit; and "why" to causal networks). However, all goals, not just linguistic queries, impose structured constraints on their solutions. The goal of getting a sloth down from a fire pole imposes different constraints than the goal of getting him up it. All that is essential for a goal to be productive is that it contains enough information to delimit the possibilities for a solution.

Conversely, not every grammatical statement that takes the form of a goal provides informative constraints. "Get the imaginary sloth down the imaginary firepole" is a productive goal; "Get the imaginary sloth down all imaginary things" is not. The issue is not the degree to which either is attainable (neither is); the issue is that the first has a defined problem space with clear starting and end states and the second does not (Simon & Newell, 1971). Similarly, "Make sure the (imaginary) sloth doesn't eat too many lollipops" is a productive goal but "Explain why (real) sloths eat so many lollipops" is not. We can think and make plans in imaginary problem spaces; we cannot think at all if the problem space fails to exist either in reality or in our imagination. People may entertain any manner of questionable goals, but we suggest that we rarely if ever consider goals that are so ill-posed that they fail to constrain the search for solutions.

Combine the idea that goals are productive and informative for thinking and planning with our ability to set goals flexibly and the result is a process that allows ideas to flourish, un-checked. The "un-checked" part is key: As we discuss below, nothing in this process guarantees that any particular idea will be beneficial now or ever. Nonetheless, insofar as human goals are both flexible and productive, they allow us to think of plans and ideas we might never have thought of had we not been trying to solve the (potentially ridiculous and insoluble) problems we were trying to solve. Our ability to set up arbitrary problems and the fact that problems provide constraints on their own solutions, accounts, we suggest, for much of the generativity of human thought.

6.5 Value to the individual

We have argued that some aspects of human goal-directed behavior are distinctive. Humans pursue goals that unfold over a lifetime or even over generations; invest energy in arbitrary goals and in goals known to be based on imaginary premises; and positively evaluate conjectures that satisfy the abstract constraints of a query or a goal even if they are knowably false. What motivates individuals to engage in these kinds of behaviors?

6.5.1 Thinking as a net cost or reward?

Most accounts of cognition treat attention, memory, processing speed, and capacity for decision-making as limited resources. Work on resource rationality has suggested that we can build more accurate models of behavior insofar as we take these costs of thinking and planning into account (Gershman et al., 2015; Lieder & Griffiths, 2020). We endorse this perspective.

However, in many contexts, thinking may still be experienced as net positive. Given the many contexts in which thinking and planning do pay off directly, we may come to assign value to thinking and planning by their frequent association with real world gains (for related ideas on "learning industriousness", see Eisenberger, 1992; Inzlicht, Shenhav, & Olivola, 2018). Thus, humans may find the activity of thinking and planning rewarding, just as we may sometimes find it rewarding to deploy physical effort (Inzlicht et al., 2018). And although individuals vary in the degree to which they value cognitive activities per se (Cacioppo & Petty, 1982), humans as a whole may value these activities more than any other species. Additionally, as noted, our capacity for long-range planning implies that we can assign value to states intermediary to a goal, even if the goal is very distant (and perhaps unachievable). That is, we experience reward for steps en route to a goal. Insofar as thinking and planning are subgoals of almost any goal, they may be intrinsically motivating, even if the functional value of the goal itself (or the probability of ever achieving it) is in doubt.

This does not mean we will find thinking about any problem whatsoever attractive. As discussed, individuals differ in their idiosyncratic interests in specific content domains. Additionally, decades of work on curiosity suggests that learners are drawn to problems of intermediate complexity: ones that are neither too easy nor too difficult to solve (Gottlieb et al., 2013; Kidd et al., 2012; Loewenstein, 1994). Something comparable may apply to our attraction to goals broadly. Some goals or problems may have too little structure to scaffold thought. Others may be too complex to readily support planning and hypothesis generation. Humans may be especially attracted to problems and goals that contain sufficient information and structure to sustain thinking and planning.

One reason to suppose we might be attuned to the degree to which we are engaged in thinking and planning is that humans have metacognitive awareness from early in childhood (Goupil & Kouider, 2016, 2019; Marazita & Merriman, 2004). Although we can fail to anticipate the pleasures associated with "just thinking" (Hatano, Ogulmus, Shigemasu, & Murayama, 2022), humans are attuned to their current state of flow or engagement (Csikszentmihalyi & Csikzentmihaly, 1990) and persist in activities that sustain it. (see Chater & Loewenstein, 2016; Danckert & Elpidorou, 2023; Lin & Westgate, n.d.; T. D. Wilson et al., 2014, for related proposals). Insofar as we perceive some kinds of cognitive effort as rewarding (Inzlicht et al., 2018), this may motivate us not only to work on problems where other payoffs are remote or unlikely; it may also motivate us to make up new problems of our own. Early researchers in intrinsic motivation discussed the possibility of autotelic behavior: behavior whose reward was the opportunity to continue engaging in the behavior itself (Csikszentmihalyi & Csikzentmihaly, 1990; Deci & Ryan, 1980; Klinger, 1969). If, as we have suggested, problems contain structured information that supports the generation of new thoughts and plans – and if thinking and planning are themselves rewarding – then making up new problems may be a means of generating new sources of intrinsic reward. That is, goals (even ones with no other apparent value in themselves) may set up a virtuous cycle where the structure of the goal both enables and motivates us to think.

In these respects, our willingness to pursue costly arbitrary goals in the absence of instrumental ends bears some resemblance to epistemic curiosity, in which we seek information for its own sake. And indeed, humans (and many other animals) are curious about information even when it is seemingly irrelevant to functional ends

(Charpentier, Bromberg-Martin, & Sharot, 2018; Vasconcelos et al., 2015) and even when it is costly to accrue (even when, for instance, it is accompanied by an electric shock Hsee & Ruan, 2016). However, we are especially motivated to seek information when it is perceived as instrumentally useful (Dubey, Griffiths, & Lombrozo, 2022; Liquin & Lombrozo, 2020b) and easy to obtain (Dan, Leshkowitz, & Hassin, 2020; FitzGibbon et al., 2020) - and information typically does have real world value in supporting our understanding of the world (Liquin & Lombrozo, 2020a). In contrast, humans can pursue goals even when doing so does not lead to information gain. Thus, we suggest the reward of pursuing seemingly arbitrary or foolish ends is not in learning or achievement but in thinking itself.

6.6 Value to Society

We turn now to the tension at the core of this article: We have set out to ask about the value of activities whose prima facie value is not obvious. If individuals experience intrinsic reward for thinking about things that support neither real world learning nor useful achievement, we are arguably engaged in a cognitive Ponzi scheme. What is the ultimate instrumental value in otherwise seemingly "useless" behavior?

6.6.1 An evolutionary argument

Here is the argument in a nutshell: Although the pursuit of some goals may look foolish, we experience real intrinsic rewards in engaging with them. Because goals contain structured information that provides valuable constraints on thinking and planning, this process serves as an engine of variation, ensuring the proliferation of new ideas and plans. Because ideas can be decoupled from the goals that motivated them, and because we are a cultural species, motivated to share novel problems and plans, these ideas are passed on intergenerationally and continue to change and adapt. Many - even most - of the ideas that result may never lead to new learning or functional outcomes. However, it may suffice that a few of them do. Our argument is necessarily speculative, but we suggest that a few important innovations - new ideas or plans we might not have come up with had we merely been trying to reduce uncertainty or achieve immediate instrumental ends - might confer large benefits to the social group, culture, or species, justifying the endeavor as a whole.

6.6.2 Cultural transmission of problems we make up for fun

One of the most salient aspects of our species is our ability to learn from one another (C. M. Wu, Vélez, & Cushman, 2022). We are a communicative species, we accumulate knowledge collectively, transmit it broadly, and pass it onto subsequent generations (Henrich, 2016; Tomasello & Rakoczy, 2003). Humans invest vast amounts of resources engaging in arbitrary problems (e.g., video games are played by 3.24 billion people, or roughly 40% of the world's population, Investopedia.com); however, we also find it rewarding to enjoy the fruits of others' labors. We watch movies and plays, listen to stories, and attend spectator sports (and make a spectator sport even of watching others play video games). We also tend to pass on the invented problems of previous generations. Ancient China and Rome are long gone; Go and Hercules endure. We suggest that our engagement with arbitrary goals and plans - our own and others - allows some of the ideas and plans we generate in pursuit of them to endure and transform.

6.6.3 Decoupling of thoughts and plans from the goals that motivated them

Critically, ideas and plans can be separated from the goals that gave rise to them and be repurposed for other ends. We owe some ideas in analytic logic to medieval monks' quest to establish incontrovertible proof of the existence of God (Glymour, 2015), and we owe some of the science of cryptography to misguided attempts to show that Francis Bacon was the author of Shakespeare's plays (Fagone, 2017). Recently, the goal of rendering graphics for animated movies and video games has inspired ideas about vision and intuitive physics (cf: "the game engine in our head"; Ullman, Spelke, Battaglia, & Tenenbaum, 2017). Even the attempt to park in every spot in a grocery store could, in principle, have led to the development of a never-before-discovered search algorithm. In this case, that failed to happen but the potential for real discoveries from foolish ends remains.

If even just a few ideas prove valuable in far-reaching ways - in genuine learning, or actionable plans that achieve previously unimaginable ends whose benefits spread widely - the investment and engagement in innumerable ends of no apparent value could be justified. Pursuing epistemic goals may not be the only - or best way to learn. The world is full of unknown unknowns. If we only engaged in learning where we expected to gain new information, we would miss the chance to learn the unexpected. Thus, even when our goals are unrealizable or foolish, they may lead to ideas of enduring practical utility, to the cultural group or humanity at large if not the individual.

Make no mistake: The flexibility and productivity of our goals, and the possibility of decoupling value from immediately useful consequences can also lead us astray: Years can be wasted looking for fountains of youth, counting angels on pins, and pursuing far more dangerous and destructive delusions. Still, the satisfaction that humans get merely from being able to keep thinking about a problem may also motivate us to work on ideas we will never see confirmed and towards ends we will never see fulfilled. Many of these ideas will be fruitless or foolish but some few may change the world.

6.7 Concluding Remarks

We have argued that human cognition is unique, not just for the sophistication of our representational abilities, but for the flexibility, structure, and value of our motivational system. We willingly incur otherwise unnecessary costs to achieve arbitrary rewards consistent with goals, problems, and constraints of our own making. This structure is productive, serving not merely to motivate thinking and planning but to enable it. The problems we construct contain information that delineates the space of their possible resolutions, supporting the search for solutions. Assigning value to arbitrary ends confers utility on sub-goals towards those ends, such that arguably costly activities like thinking and planning can have positive utility, even if the ideas and plans have no other consequences for prediction or action. Our ability and willingness to think about whatever we like, liberated from the demands of functional outcomes may pay off for the species in our ability to think of ideas we would not have thought of had we not been trying to achieve a particular, idiosyncratic goal. The ideas we generate can then be detached from the goals that originally inspired them, shared with others, and put to new, and sometimes transformative, ends.

Outstanding Questions

- What challenges does the flexibility of human goals pose for our theories of exploration, planning, and inference? How do we translate abstract goals ("rescue the sloth") into specific reward functions or satisfaction criteria that support planning and hypothesis generation? How can we leverage computational models of planning and inference to explain how humans represent, select among, and eventually achieve arbitrary goals?
- How do we represent and search within the space of possible goals? Is the space of possible goals a uniform distribution over any possible desire, or might they be organized and structured - just as our beliefs may be organized around concepts and hierarchically structured in the form of intuitive theories?
- How do we use the information contained in a problem to solve it?
- What does it mean for a goal to set up a well-specified problem? How do we predict expected cognitive engagement for a novel goal? To what extent do individuals converge when evaluating a goal, and how might these judgments vary with respect to differences in individual knowledge, preferences, and abilities?
- What combination of domain knowledge, individual preferences and abilities, and contingent opportunities leads us to take on the array of goals we do? Where do our idiosyncratic goals themselves come from? Do we develop expertise in creating and selecting well-posed problems? If so, is that domain-specific or domain-general and how is it acquired?

Figure 1: Children's behavior is conditionally rational with respect to their goals

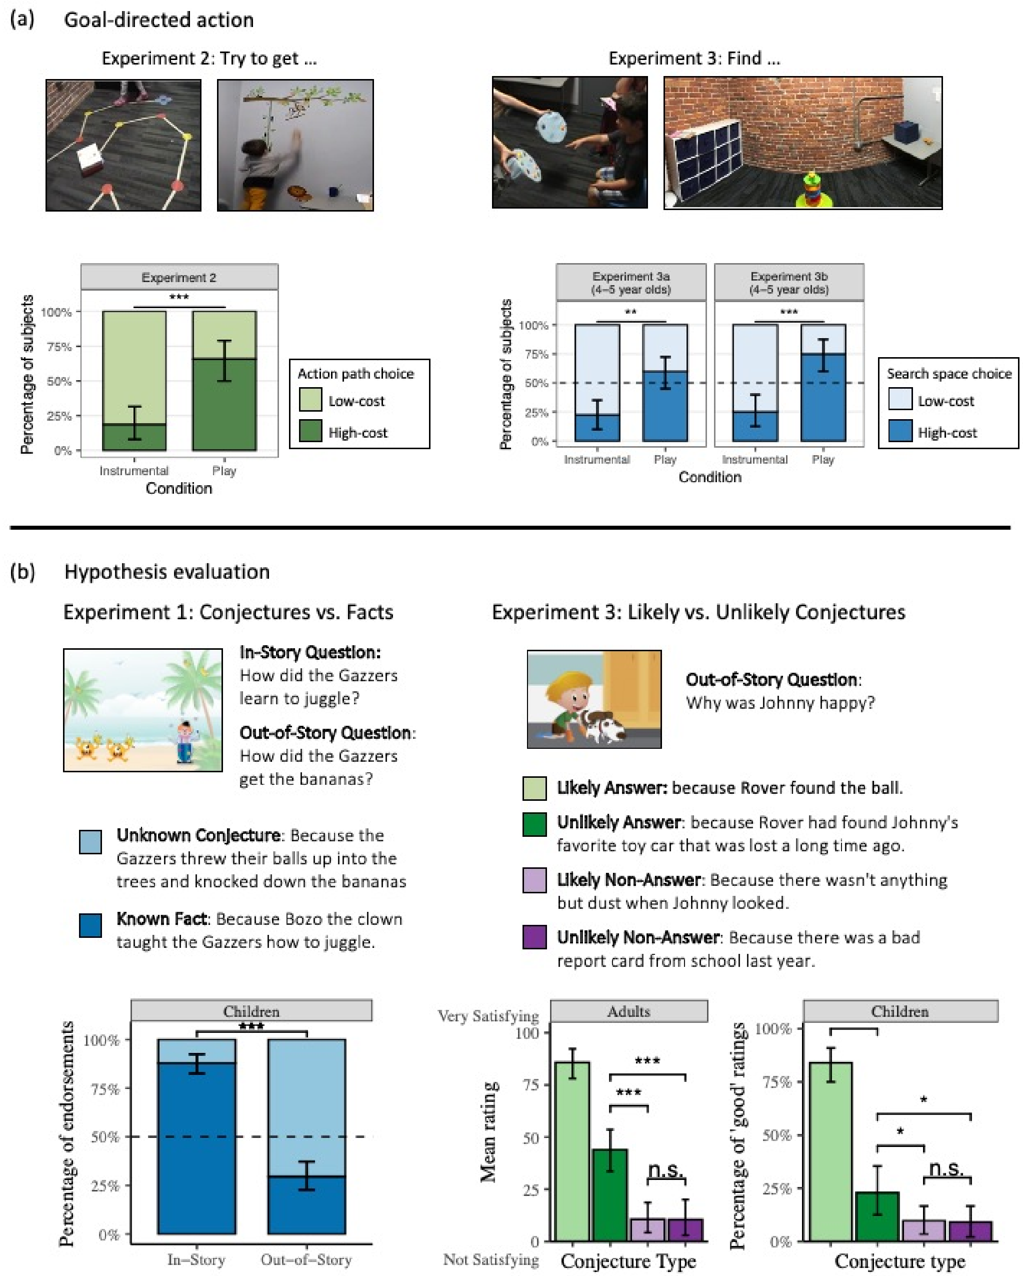

Figure 6-1: Children's behavior is conditionally rational with respect to their goals. We have argued that people voluntarily incur unnecessary costs for arbitrary rewards, but that goals, even arbitrary ones, provide valuable constraints that constrain thinking and planning. In (A), we show the results of two experiments (Chu & Schulz, 2021, see Chapter 4 for details) in which children chose more costly actions when told to "Go play" and obtain a target than when simply asked to "Go get" the target. Nonetheless, children's actions were conditionally rational: children behaved efficiently (e.g., sticking close to the path, jumping directly up to the pencil, checking adjacent locations) with respect to their self-imposed goals. In (B) we show the results from two experiments in which participants were asked to answer questions about a story (Chu & Schulz, in press, Chapter 3). In Experiment 1, we showed that despite children's well-established preference for reliable evidence (Harris et al., 2018; Koenig et al., 2019), children preferentially endorsed speculative conjectures that could achieve the goal of answering the question when known facts could not. In Experiment 3, we showed that both adults and children endorse even highly improbable conjectures that could in principle answer the question over probable (and improbable) answers that could not. Thus, we evaluate proposals with respect to how well they achieve the goal of potentially answering the question, not merely with respect to the facts.

Chapter 7

Afterword

> For me, the biggest mystery of all is still the question I posed at the beginning: why is problem-solving sometimes so fun? That is, why do human beings go looking for problems to solve instead of just doing things we can already do? What function does problem-solving serve to make it so attractive and to justify its fundamental role in childhood activities? In a real sense, problem-solving is at the heart of what we mean by intelligence. The ability to identify a goal, work out how to achieve it, and carry out that plan is the essence of every intelligent activity.
>
> Stephanie Thornton, 1995, Children Solving Problems, p. 126

This thesis introduced the idea that human cognition is organized around goal representations, which constrain thought by providing a source of *information* for reasoning in novel situations, as well as a source of *value* in decision-making. In support of this idea, I presented empirical evidence from three case studies. First, our goal of answering a question can lead us to entertain novel conjectures even when they are unfamiliar, lack evidential support, or are a priori low in plausibility (Chapter 3). Second, in play, freed from more instrumental concerns, we explore new ways of achieving the same goals by adopting unnecesary subgoals and constraints on our plans (Chapter 4). And by simply choosing a goal, we have constrained what our future selves will value: we assess choices and actions with respect to our chosen goals, rejecting lower-cost alternatives that we would have just a moment ago been quick to select (Chapter 5). Empirically, this thesis contributes evidence that adults and young children represent and use their goals to judge new ideas, to make new plans, and to choose how to deploy their effort.

Theoretically, this thesis contributes new questions and directions for the study of flexible human cognition. Each empirical study demonstrated behaviors that pose challenges for existing accounts of rational inference and decision-making; participants seemed to go beyond simple definitions of accuracy and efficiency when evaluating proposals or taking actions. Instead, these results suggest an extraordinary amount of flexibility in how we reason and pursue our goals. This flexibility is at once distinctive and characteristic of higher order human cognition, as manifest in our capacity for creating and solving new problems throughout the lifespan, and also distinctive and characteristic of early childhood, as manifest in play (Chapter 2).

Furthermore, beyond simply *affecting* how we think, the goals we adopt might be the very structures that enable us to think at all. When prior knowledge and data are lacking - for example in situations where the space of possibilities is large, complex, or highly uncertain - goal representations may offer a critical guide for how we reason or plan. By generating and pursuing new goals, we may discover new ideas and plans we would not have explored otherwise, had we been focused on a singular objective such as *maximize learning*. And by finding it rewarding to generate new goals and plans, independent of the expected outcomes, we have a self-sustaining engine of variation that may contribute to our species' distinctively powerful ability to innovate and adapt in novel situations (Chapter 6).

We can now return to the central question of this dissertation: *What kind of mind yields the rich diversity of goals that humans entertain and adopt, as well as the flexibility with which we reason and plan in pursuit of these goals – not to mention the fun we get out of doing so?* The answer offered in this dissertation is two-fold: such a mind is a (1) ***pragmaticist*** in that cares deeply about its goals, and assesses inferences and plans with respect to those goals (Dewey, 1916; James, 1907; Peirce, 1878); it is also (2) a ***playful*** mind that takes pleasure in the process of generating and pursuing novel goals, independent from expected outcomes and action costs.

While these chapters capture important aspects of the human ability to engage in ad hoc, goal-directed reasoning, there are many open questions about their specific cognitive, computational, and developmental basis (e.g., Section 6.7). What other challenges does the flexibility of human goals pose for our theories of exploration, planning, and inference, and what empirical data will distinguish those theories? How can we leverage computational models of planning and inference to explain how humans represent, select among, and eventually achieve arbitrary goals? In the remainder of this chapter, I describe some specific directions for future research.

Given a goal, how do we generate and assess plans on the fly?

In unfamiliar situations without relevant prior knowledge and evidence, holding arbitrary goals without obvious utilities, *how* do we generate candidate hypotheses and plans, and ignore or reject those which appear irrelevant? It is a wide open question how we identify and bring to bear *just* the information that is relevant for any given situation. In this sense, we are after a solution to the Frame Problem (Dennett, 1984)[1]. Prior work has demonstrated the potential of modeling reasoning as stochastic search in a space of possibilities (e.g., Ullman, Goodman, & Tenenbaum, 2012). This sampling may be constrained by not just what is possible or probable, but also by what is valuable (J. Phillips, Morris, & Cushman, 2019). However, we do not yet have an account of how people might compute the value of different options across domains, especially with respect to novel goals.

In computational terms, this puzzle might be framed as the problem of how to *conditionalize* inferences on our goals. An open question then is: how we are representing the goals in a way that can flexibly enter so many different kinds of reasoning, from evaluating to generating possibilities? One promising candidate from computational theory may be representations of *types* (Sosa & Ullman, 2022).

How is planning both costly and rewarding?

Resource rational models went beyond previous accounts of rationality by expanding our notion of costs, to include the cost of thinking. However, we might consider expanding our notions of rewards to include something about our goals. As discussed in Chapters 4 and 6, one specific possibility is that we are sensitive to - and seek out - opportunities for thinking and cognitive engagement. Perhaps we are sensitive to how many potential solutions exist, how sparse or difficult to find a single solution, or how clustered this hypothesis space is. These factors might all predict the amount and kind of cognitive computation we might expend. Might we see a monotonic relationship

[1]Thanks to Lionel Wong, Alex Lew, and Tyler Brooke-Wilson for introducing this link.

between expected computations and motivation, or perhaps a U-shaped relationship similar to research on information complexity and curiosity (e.g., Kidd et al., 2012)?

In fact, the kinds of *opportunities for thinking* we are sensitive to may be as diverse as the kinds of thinking we are able to engage in. Our assessment and preferences about planning problems may reflect individual differences in our cognitive processes, including those resulting from maturation or learning. For example, different planning puzzles (e.g., Rush Hour, block construction, or flood-filling tiles) might be formally equivalent when represented as graphs, but we might think of one puzzle as rewarding and another as frustrating, depending on our representational and inferential resources in the domain, or our ability to convert a puzzle to whichever representational format our solution algorithms work in.

Further, these assessments may depend not only about our mental resources, but also our meta-cognitive beliefs about our mental computations. Recent work suggests that adults can represent not just others' mental states, but also their cognitive processes, such as when we correctly recover someone's goal from inefficient or even unsuccessful attempts by accounting for resource-limited planning (Alanqary et al., 2021), or when we infer that someone who thinks slowly on an easy problem is probably daydreaming (Berke & Jara-Ettinger, 2021). Extending this work, I expect that reasoning about the *utility* that others place on different planning processes may help us to negotiate goals in social contexts, such as when delegating work, choosing what games to play together, or designing classroom projects to foster learning in self-directed ways. This hypothesis aligns with recent studies showing that children choose tasks for themselves and others with respect to the agents' goals and abilities. For example, children choose harder tasks when trying to have fun than when trying to achieve the goal (Chapter 4) or trying to win (Goddu, Rule, Bonawitz, Gopnik, & Ullman, 2022), and children will also distribute tasks to others depending on their relative capabilities (Baer & Odic, 2022; Magid et al., 2018).

Linking planning to emotion and motivation

If thinking can be both rewarding and costly, might our subjectively felt emotions and motivations be related to our ongoing or expected cognitive processes? For example, might our feelings of excitement or boredom, or of frustration or enthusiasm, derive from computations over the expected kind of thinking that a problem will engage? Some recent work suggests that both adults and young children are able to represent and make inferences about an agent's cognitive processes, such as from how long they take to respond. When we see someone solve a puzzle extremely quickly, we may infer that they have seen the puzzle before (as opposed to solving it for the first time), or that they cheated, or that the puzzle is actually much easier than we thought (Berke & Jara-Ettinger, 2021; Richardson & Keil, 2022). And while past work has looked at how emotions, such as surprise and confidence, can reflect epistemic states in ourselves and others (Baer, Gill, & Odic, 2018; Y. Wu & Gweon, 2019), future work might investigate "metacognitive feelings" that reflect cognitive *processes*. Future work might also investigate children's intuitive theories about cognitive processes, and how children can use expectations about task demands and agents' cognitive abilities to predict emotional reactions, or to prospectively choose tasks contingent on the particular agent and their goals.

How do we explore the space of potential goals?

Exciting recent work in robotics and machine learning suggest that one effective way to acquire open-ended skill repertoire might be to build *autotelic agents*: intrinsically motivated learning agents that can learn to represent, generate, select and solve their own problems (e.g. Colas et al., 2022; Ellis et al., 2021; Forestier et al., 2017; Laversanne-Finot, Péré, & Oudeyer, 2021; Schmidhuber, 2013). These ideas are inspired by human development, such as the intrinsically motivated play of young children. However, open questions remain about how humans engage in goal exploration. Faced with novel goals, how do we decide what to pursue, given our individual knowledge, skills, and preferences? More generally, how do people represent and explore the space of potential goals? What features do we use to assess goals beyond their expected outcome and information gain?

While intuitive, we do not yet have a general account of how humans assess how interesting or difficult a task will be. As one example, consider what is similar or different in reasoning about: how hard a research project will be, how hard it will be to build a block tower (Bennett-Pierre & Gweon, 2018; Dietz, Landay, & Gweon, 2019; Leonard, Bennett-Pierre, & Gweon, 2019; Yildirim et al., 2019), how long it will take us to establish a habitable site on Mars, or to train for the Boston Marathon? Tasks often differ in multiple dimensions, and if people represent both the costs and rewards of planning, how do those computations serve as inputs to goal exploration? One source of complexity is the fact that perceived costs may interact with perceived rewards: a seemingly difficult or uncertain task might simultaneously offer more room for skill improvement or information gain. How do we integrate these different cues? Finally, simulation and counterfactual reasoning may also play a role. One recent study found that people's judgments of how interesting a block tower would be to build correlated strongly with how precarious the tower looks, which can be computed from integrating geometric features and physical simulation (Holdaway et al., 2021). Thus, we may assess not just the final end state of a goal, but also counterfactuals about alternative paths and outcomes along the way.

Creative processes

Good ideas, as well as good problems, rarely come to our minds as neat packages ready to be deployed. What role does the *process* of generating problems and ideas play in the form they eventually take? How do we iterate over our ideas and make them better? By hypothesis, the problems and goals we have serve as metrics for assessing improvement. But in the process, might we also gain clarity on the problems themselves? If yes, this implies a positive feedback loop for innovation.

Here is a conjecture: Just as we might "learn by thinking" without any new information from the world – such as by explaining ideas to ourselves and others (Lombrozo, 2019; Walker, Lombrozo, Legare, & Gopnik, 2014; Walker, Lombrozo, Williams, Rafferty, & Gopnik, 2016), or engaging in thought experiments (Bascandziev & Carey, 2022; Bascandziev & Harris, 2019) – perhaps we also "learn by expressing" the same information in different ways. Lacking a concrete formulation, I will illustrate these idea in two examples.

One example comes from articulating scientific ideas. At one of my PhD committee meetings, I was advised to practice my presentation skills by giving more research talks. So I signed up to give four external talks, each about ten days apart. Each week, I found myself improving at delivering the intended argument, but surprisingly, the argument and intended message itself also improved with each iteration. Some of this change is surely attributable to mere practice, or new information from audience feedback. But speculatively, some of the improvement might be attributable to the iterative process of preparing these talks. Throughout the process, I had a continuous abstract goal ("deliver a good talk") that was hard to plan towards, and multiple changing concrete goals ("communicate these ideas to this particular audience" or ") which clarified what would constitute a "good talk" and offered more informative constraints on my planning. I also wonder how my ideas would have changed had I been iterating on four papers, or four diagrams. What role do external representations and modes of expression play in this process of thinking better?

Another example of the benefit of externalizing our problems comes from writing code. A close friend, who is much better at programming, once suggested this hack: if you are stuck on a coding problem, simply place a puppet on the desk, and tell it what your problem is. In the process you might discover a new insight, or a piece of knowledge or resource that you might have forgotten about. A few years of practicing this hack and I can now imagine talking to different advisors, depending on the problem. Sometimes, however, I find it much more useful to sketch with a whiteboard, or to write in longform. Why might different modes of expression work better or worse for different problems? Might we intuitively recognize that some cognitive tools are more suitable for some problems and agents than others?

Social contexts

All of the behaviors discussed in this dissertation also exist within rich social contexts. We talk about and share our goals with others, and it's often really productive and fun to work together on a problem. But communication and collaboration require rich social inferences about what our partners know, what they want, and what their abilities are – each of these may be highly underdetermined and complex inferences. Yet, I suspect children succeed in many of these inferences, if only to support playing well with others. How do those inferences work? And sometimes we get goals that others suggest or assign, and we may choose goals for other agents to pursue, such as in teaching contexts. How do we communicate and collaborate in order to play & think together? I suspect that studying the social contexts of problem solving and problem selecting might yield fruitful insights for connecting this work to educational contexts.

How do we get better at creating new problems and solutions?

Classic work in cognitive science has investigated the nature of expertise and problem representation on tasks ranging from physics problems to chess and tetris (Anderson, 1982; Chi, Feltovich, & Glaser, 1981; Gray & Banerjee, 2021). How might we extend these theories to account for the expertise in not just solving, but also creating and designing new goals and problems? Relatedly, what is the nature of expertise involved in designing not just informative learning experiences but also engaging thinking experiences? How domain-specific is this expertise?

Some professions capitalize on an ability to identify – or to modify or even create – problems that are fun and engaging (e.g., game designers and TV show writers). Are good problem setters also good problem solvers, and how might we acquire this sort of expertise? Perhaps it's time to visit the playground again.

Conclusion

This work has attempted to take seriously our ability and inherent motivation to invent and pursue novel goals. I have argued that doing so will lead us towards asking new questions about the rich complexity of human goal-directed behavior, and a better understanding of the nature of flexible human intelligence and its origins in early childhood. In the process, we may find in children's play the very seeds of human innovation and creative problem solving.

Appendix A

Supporting Information for Chapter 3

Table A.1: Stimuli text for Experiment 1

Item	Text
Training Story 1	One day, Tommy wanted to visit the big castle at the end of the road. The castle was very far away, so Tommy had to ride a bike. Tommy rode his bike all the way to the Castle and he was so happy!
In-Story Question	How did Tommy get to the castle?
Fact Answer	He walked to the castle.
Conjecture Answer	He rode a bike to the castle.
Training Story 2	Inside the castle, Tommy found an ice cream store. There were lots of different ice cream flavors, but Tommy's favorite is Chocolate. So Tommy got a chocolate ice cream and ate it all up.
In-Story Question	What is Tommy's favorite ice cream?
Fact Answer	Chocolate
Conjecture Answer	Strawberry
Test Story 1	Here are some small Daxes and here are some big Blickets. The big Blickets are hat makers and they made hats for the small Daxes.
In-Story Question	Why are the small Daxes wearing hats?

Continued on next page

Table A.1: Stimuli text for Experiment 1 (Continued)

Out-of-Story Question	Why are the Blickets bigger than the Daxes?
Fact Answer	Because the big Blickets made the hats for them .
Conjecture Answer	Because the Blickets are older than the Daxes .
Test Story 2	Here are some spotted Wugs and here are some furry Feps. The spotted Wugs are allergic to the Fep's fur and keep sneezing. The furry Feps think the sneezing Wugs are funny and giggle every time the Wugs sneeze .
In-Story Question	Why are the Wugs sneezing?
Out-of-Story Question	Why are the Feps furry?
Fact Answer	Because theWugs are allergic to the Feps fur .
Conjecture Answer	Because Feps go up to the mountains and their fur keeps the Feps warm .
Test Story 3	Here are some juggling Gazzers. A clown named Bozo taught them to juggle. Juggling Gazzers love to eat bananas. But the bananas all grow at the top of very tall trees and the Gazzers can't climb trees. But here the Gazzers are! Eating bananas .
In-Story Question	How did the banana eating Gazzers learn to juggle?
Out-of-Story Question	How did the Gazzers get the bananas?
Fact Answer	Because Bozo the clown taught the banana eating Gazzers how to juggle .
Conjecture Answer	Because the Gazzers threw their balls up into the trees and knocked down the bananas .
Test Story 4	Here are some Duffs who need haircuts. Their hair got so long that they couldn't see where they were going and they tripped and dropped their favorite toy down a deep hole. Poor Duffs! But look they managed to rescue their toy!
In-Story Question	How did the Duffs' toy fall down the deep hole?
Out-of-Story Question	How did the Duffs rescue their toy?
Fact Answer	Because the Duffs' hair was in their eyes and they couldn't see and they tripped and dropped their toy .

Continued on next page

Table A.1: Stimuli text for Experiment 1 (Continued)

Conjecture Answer	Because the Duffs tied their long hair into a rope and made a ladder with it and used it to climb down and get the toy .

Note: With one exception, Elmo always provided the Fact answer and Cookie Monster always provided the Conjecture answer. One exception is on Training Story 2, in which Elmo provided the incorrect Conjecture answer and Cookie Monster provided the correct Fact answer. This change ensured that the training trials showed both puppets paying attention to information provided in the story.

Table A.2: Stimuli text for Experiment 3

Item	**Text**
Training Story 1	This is Tina. Tina is having breakfast. She's eating pancakes. After breakfast, she went outside to play. When she was done, it was time for lunch. Tina comes back into the kitchen. Her brother comes into the room and asks, "Hey Tina, what did you have for breakfast?"
Question	What did Tina have for breakfast?
Likely Answer (Fact)	Tina ate some pancakes for breakfast.
Likely Non-Answer	Tina played in the tree-house for breakfast.
Training Story 2	This is Tommy. Tommy went to the ice cream shop. He bought his favorite ice cream and ate it all up! (only an empty cone was shown)
Question	What did Tommy get at the ice cream shop?
Unlikely Non-Answer	Tommy got tomato soup.
Likely Answer	Tommy got chocolate ice cream.
Test Story 1	Everyday, Jessie and her brother Tom take the bus home from school together. But today, Jessie is home sick. Jessie stayed upstairs in bed all day. It's the afternoon and her brother usually comes home now. She hears the dog start to bark. The dog is looking out the window, barking loudly!

Continued on next page

Table A.2: Stimuli text for Experiment 3 (Continued)

Question	Why was the dog barking?
Likely Answer	Because her brother Tom just got off the bus, and he was walking to the front door.
Unlikely Answer	Not because it was her brother Tom, but, because it was a cat in the front yard and it had tipped over the garbage can.
Likely Non-Answer	Not because it was her brother Tom, but, because there was a flower in the front yard.
Unlikely Non-Answer	Not because it was her brother Tom, but, because there were dirty dishes in the kitchen sink.
Test Story 2	This is Billy. Billy always gets mad at his naughty brother Sam. Sam always takes Billy's toys from his box, and leaves them lying around his room. So Billy decides to walk over to his brother's room. Billy opened his brother's door. When he looked inside, he started yelling.
Question	Why was Billy yelling?
Likely Answer	Because Billy's toys were lying all around. Billy got so mad he yelled and started waving his arms angrily at his brother.
Unlikely Answer	Not because Billy's toys were lying all around, but, because when he opened the door, the lamp had tipped over and the curtains were on fire!
Likely Non-Answer	Not because Billy's toys were lying all around, but, because when he opened the door, the curtains were blowing in the wind.
Unlikely Non-Answer	Not because Billy's toys were lying all around, but, because when he opened the door, his brother was doing homework.
Test Story 3	This is Mary. Mary wore a brand new skirt out to play. But she ripped the new skirt on the way home. She came home feeling very upset and crying really hard.
Question	Why was Mary crying?
Likely Answer	Because Mary really liked her beautiful flowered skirt and now it is ripped and completely ruined.
Unlikely Answer	Not because it was Mary's beautiful flowered skirt, but, because it was her sister's skirt and Mary's afraid her sister will get really mad at her.

Continued on next page

Table A.2: Stimuli text for Experiment 3 (Continued)

Likely Non-Answer	Not because it was Mary's beautiful flowered skirt, but, because her friend made a funny joke earlier that day.
Unlikely Non-Answer	Not because it was Mary's beautiful flowered skirt, but, because it was an old skirt that she didn't like very much and already had rips and tears.
Test Story 4	This is Johnny. Johnny loves to play fetch with his dog Rover. Sometimes Johnny throws the ball too far, and it rolls into the shed. Rover runs after the ball. Rover looks under the cabinets and starts barking and wagging his tail. Johnny runs in and looks under the cabinet. He then pats his dog and says excitedly, "Good boy!"
Question	Why was Johnny happy?
Likely Answer	Because Rover found the ball. He followed the green tennis ball all the way into the shed. Johnny thought Rover was so clever.
Unlikely Answer	Not because Rover found the ball, but, because Rover had found Johnny's favorite toy car that he had lost a long time ago.
Likely Non-Answer	Not because Rover found the ball, but, because there wasn't anything but dust when Johnny looked.
Unlikely Non-Answer	Not because Rover found the ball, but, because there was a bad report card from school last year.
Test Story 5	This is Sally. Sally was looking forward to her best friend's birthday party. Her best friend had just mailed out the invitations, and Sally was hoping to get one soon. Sally walked to her mailbox and saw a shiny white envelope. Sally opened the envelope and jumped up and down excitedly when she read it!
Question	Why was Sally so excited?
Likely Answer	Because she got invited to her best friend's birthday party. She was so excited that she couldn't stop jumping up and down.
Unlikely Answer	Not because it was a birthday invitation, but because it was a letter from school saying she won the story competition.
Likely Non-Answer	Not because it was a birthday invitation, but because it was a notice from the library saying she forgot to return her books.

Continued on next page

Table A.2: Stimuli text for Experiment 3 (Continued)

Unlikely Non-Answer	Not because it was a birthday invitation, but because it was a note from her teacher saying that she had to do extra work after school.
Test Story 6	Every Saturday afternoon Penny rides her bike to go to the beach and play with her friends. She has lots of friends and they always play happily. Today, there is construction work between Penny's house and the beach. Penny isn't at the beach yet. Penny's friends wonder where she is.
Question	Why is Penny late?
Likely Answer	Because Penny was still trying to find a way to get around the construction. There was a big yellow truck blocking her bike path.
Unlikely Answer	Not because Penny was trying to get around the construction, but, because she has chickenpox and has to stay home in bed.
Likely Non-Answer	Not because Penny was trying to get around the construction, but, because it was a gorgeous sunny day at the beach, perfect for swimming.
Unlikely Non-Answer	Not because Penny was trying to get around the construction, but, because the construction trucks were really quiet.

Note. Participants were excluded for incorrectly responding to either training trial.

A.1 Experiment 1 analysis by "why" and "how" questions

In Experiment 1 we used both causal "why" questions and problem-solving "how" questions. While this contrast was designed to cover a broad range of explanatory questions, random effect coefficients from the main analysis reported in the manuscript suggested that participants were more likely to endorse the corresponding answers on "how" questions than "why" questions. Thus, we conducted a post-hoc analysis to examine potential effects of question prompt and the interaction between question prompt and question type on whether participants chose the appropriate explanation for that question type (0=Inappropriate, 1=Appropriate). We constructed a logistic mixed-effects model with question prompt (0=why, 1=how) and question type (0=In-Story, 1=Out-of-Story) as fixed effects, and random intercepts for subject.

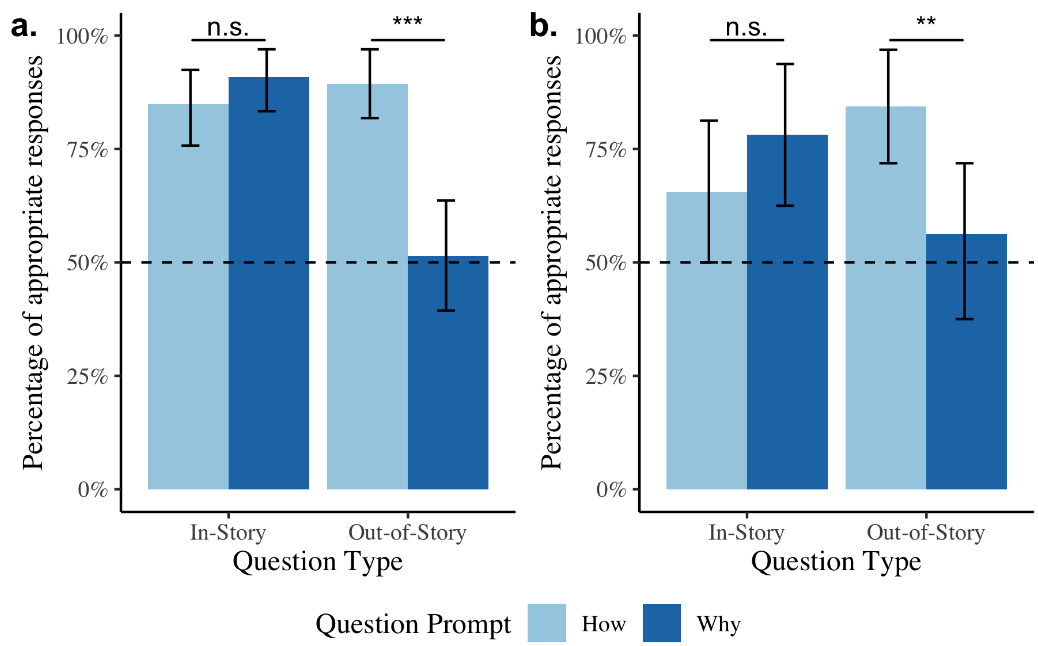

Figure A-1: Children's ratings (averaged across two test trials for each question type) in Experiments 1a (N=66, mean: 6.04 years; range: 4.00-7.93) and 1b (N=32; mean: 5.03 years; range: 4.15 – 5.92). Children appropriately preferred facts for questions that could be answered in the story, regardless of question prompt. Children preferred conjectures when the question could not be answered in the story, but only for "How" questions and not for "Why" questions. Paired t-test, $^{**}p < .01$, $^{***}p < .001$

In Experiment 1a (N=66, mean: 6.04 years; range: 4.00-7.93), we found a significant question prompt by question type interaction (β= -2.65, z = -3.67, $p < .001$). To explore these interactions, we conducted Tukey-adjusted pairwise comparisons on the proportion of appropriate answers chosen for In-Story versus Out-of-Story questions using either prompt. While children appropriately chose Fact explanations in responses to In-Story questions with both *How* and *Why* prompts, and appropriately chose Conjecture explanations to Out-of-Story questions with How prompts, they were significantly less likely to choose Conjecture explanations for Out-of-Story questions with Why prompts (see figure S1a).

We conducted the same analysis with data from Experiment 1b (N=32, mean: 5.03 years; range: 4.15 – 5.92) and found a significant question prompt by question type interaction effect (β= -2.62, z = -2.67, p = .007). Using Tukey-adjusted pairwise comparisons, we again found that children chose appropriate explanations for both In-Story questions and for Out-of-Story questions with *How* prompts, but were significantly less likely to respond appropriately for Out-of-Story questions using the *Why* prompt.

A.2 Experiment 4 analysis by age

In the main text we report a significant age by conjecture type interaction when age was coded categorically (younger 4-5-year-olds vs. older 6-7-year-olds). Here we replicate the analysis using age as a continuous variable (mean-centered).

We built a logistic regression model with the fixed effects of conjecture type, age, and an age by conjecture type interaction, with age in years. This model explained more variance than either the conjecture type only model ($\chi^2(3) = 19.25$, $p < .001$) or the model with conjecture type and age but no interaction ($\chi^2(2) = 19.15$, $p < .001$). Inspection of estimated marginal slopes indicated that with increasing age, children became more likely to endorse Unlikely Answers ($\beta = 0.48$), but less likely to endorse conjectures that did not answer the question (Likely Non-Answers: $\beta =$ -0.85); Unlikely Non-Answers: $\beta =$ -0.84).

Appendix B

Supporting Information for Chapter 4

In the main manuscript, Experiment 3b, we reported an exploratory analysis of the effect of age on children's selection of the smaller or larger search space (Figure 4-5, main manuscript). Here we conduct an exploratory reanalysis of Experiment 3b, asking if the effect of age and condition might differ across both tasks (Boxes and Buttons).

We fit a logistic mixed effects regression model with condition (Instrumental or Play), age (in months), and task (Boxes or Buttons) as predictors. We included all two- and three-way interactions, as well as a random by-subject intercept. Note that compared to the original model reported in the manuscript, which excluded the task variable, this expanded model did not explain any significant additional variance ($chi^2(4) = 2.40$, $p = 0.66$).

Investigating simple slopes of age within each task and condition combination, we found that in the Play conditions, children of all ages preferred to choose the high-cost larger search space (see Figure B-1). There was no age effect on the Play conditions for either task (Boxes: $\beta = .019$, OR = 1.02, 95% CI:0.92-1.13, $p = .71$; Buttons: $\beta = .014$, OR = 1.01, 95% CI:0.94-1.10, $p = .74$).

In contrast, in the Instrumental condition, older children were increasingly likely to choose the low-cost smaller search space in both tasks, although these estimates did not differ significantly from zero (Boxes: $\beta = -.083$, OR = 0.92, 95% CI:0.84-1.01, $p = .08$; Buttons: $\beta = -.07$, OR = 0.93, 95% CI: 0.85-1.02, $p = .14$).

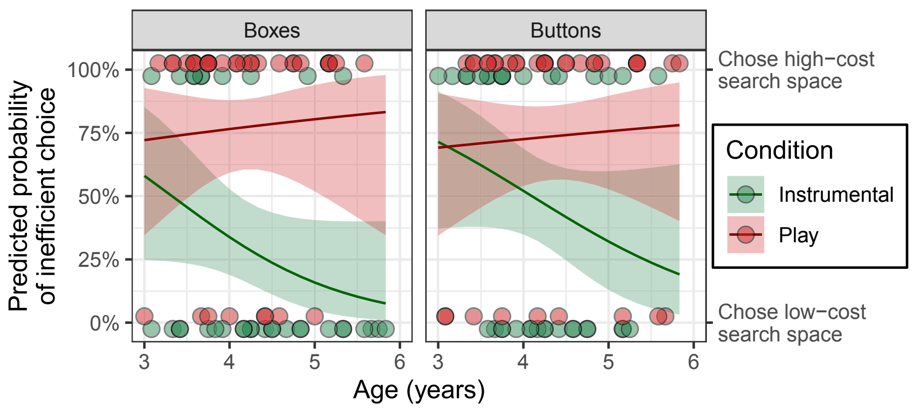

Figure B-1: Age differences in search efficiency in Experiment 3b (N=69). In both tasks, older children were increasingly likely to make low-cost, efficient choices in the Instrumental search task (green line). However, in play, children of all ages preferred the larger search space. Each circle represents the choice of one participant and lines show predicted probability of making each choice across the ages tested; shaded regions indicate 95% confidence intervals on best-fit logistic regression estimates.

References

Aboody, R., Zhou, C., & Jara-Ettinger, J. (2021). In Pursuit of Knowledge: Preschoolers Expect Agents to Weigh Information Gain and Information Cost When Deciding Whether to Explore. *Child Development*, *92*(5), 1919–1931. doi: 10.1111/cdev.13557

Agrawal, P., Nair, A., Abbeel, P., Malik, J., & Levine, S. (2016). Learning to poke by poking: Experiential learning of intuitive physics. *arXiv*, *1606.07419*.

Alanqary, A., Lin, G. Z., Le, J., Zhi-Xuan, T., Mansinghka, V. K., & Tenenbaum, J. B. (2021). Modeling the mistakes of boundedly rational agents within a bayesian theory of mind. *CoRR*, *abs/2106.13249*. Retrieved from `https://arxiv.org/abs/2106.13249`

Alessandri, S. M. (1991). Play and social behavior in maltreated preschoolers. *Development and Psychopathology*, *3*(2), 191–206. doi: 10.1017/s0954579400000079

Alexander, J. M., Johnson, K. E., Leibham, M. E., & Kelley, K. (2008). The development of conceptual interests in young children. *Cognitive development*, *23*(2), 324–334. doi: 10/dxtbq8

Allison, M.-L., Reed, R., Michels, E., & Boogert, N. J. (2020). The drivers and functions of rock juggling in otters. *Royal Society Open Science*, *7*(5), 200141. doi: 10.1098/rsos.200141

Amso, D., & Johnson, S. P. (2006). Learning by selection: Visual search and object perception in young infants. *Developmental Psychology*, *42*(6), 1236–1245. doi: 10.1037/0012-1649.42.6.1236

Andersen, M. M. (2022). Play in predictive minds: A cognitive theory of play. *Psychological Review*. doi: 10.1037/rev0000369

Anderson, J. R. (1982). Acquisition of cognitive skill. Cités par moi. *Psychological review*, *89*(4), 396–406. doi: 10.1037/0033-295x.89.4.369

Anderson, J. R. (1990). *The Adaptive Character of Thought*. Lawrence Erlbaum Associates.

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*(3), 409–429. doi: 10.1037/0033-295x.98.3.409

Arkes, H. R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes*, *35*(1), 124–140. doi: 10/crq2tt

Arkes, H. R., & Hutzel, L. (2000). The role of probability of success estimates in the sunk cost effect. *Journal of Behavioral Decision Making*, *13*(3), 295–306.

Aslin, R. N. (2007). What's in a look? *Developmental Science*, *10*(1), 48–53.

Baer, C., Gill, I. K., & Odic, D. (2018). A domain-general sense of confidence in children. *Open Mind*, *2*(2), 86–96. doi: 10.1162/opmi_a_00020

Baer, C., & Odic, D. (2022). Mini managers: Children strategically divide cognitive labor among collaborators, but with a self-serving bias. *Child Development*, *93*(2), 437–450. doi: 10/gnbdws

Baillargeon, R. (1987). Object permanence in 31/2-and 4$1/2$-month-old infants. *Developmental Psychology*, *23*(5), 655. doi: 10.1037/0012-1649.23.5.655

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

Baldwin, J. D., & Baldwin, J. I. (1974). Exploration and social play in squirrel monkeys (saimiri). *American Zoologist*, *14*(1), 303–315. doi: 10.1093/icb/14.1.303

Bálint, V., Csibra, G., & Kovács, A. M. (2021). Infants' interpretation of information-seeking actions. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (pp. 472–478).

Baranes, A., & Oudeyer, P.-Y. (2010). Intrinsically motivated goal exploration for active motor learning in robots: A case study. In *IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 1766–1773). Ieee.

Barrett, L., Dunbar, R. I., & Dunbar, P. (1992). Environmental influences on play behaviour in immature gelada baboons. *Animal Behaviour*, *44*(1), 111–115. doi: 10.1016/s0003-3472(05)80760-2

Barto, A. G. (2013). Intrinsic Motivation and Reinforcement Learning. In G. Baldassarre & M. Mirolli (Eds.), *Intrinsically Motivated Learning in Natural and Artificial Systems* (pp. 17–47). Berlin, Heidelberg: Springer Berlin Heidelberg. doi: 10.1007/978-3-642-32375-1_2

Bascandziev, I., & Carey, S. (2022). Young children learn equally from real and thought experiments. In *Proceedings of the 44th annual meeting of the cognitive science society.*

Bascandziev, I., & Harris, P. (2014). In beauty we trust: Children prefer information from more attractive informants. *British Journal of Developmental Psychology*, *32*, 94–99. doi: 10/gk28nz

Bascandziev, I., & Harris, P. L. (2019). Can children benefit from thought experiments. *The scientific imagination*, *262*.

Bateson, P., Martin, P., & Young, M. (1981). Effects of interrupting cat mothers' lactation with bromocriptine on the subsequent play of their kittens. *Physiology & Behavior*, *27*(5), 841–845. doi: 10.1016/0031-9384(81)90051-2

Bateson, P., Mendl, M., & Feaver, J. (1990). Play in the domestic cat is enhanced by rationing of the mother during lactation. *Animal Behaviour*, *40*(3), 514–525. doi: 10.1016/s0003-3472(05)80532-9

Beckel, A. L. (1991). Wrestling play in adult river otters, lutra canadensis. *Journal of Mammalogy*, *72*(2), 386–390. doi: 10.2307/1382111

Becker, G. S. (1968). Crime and punishment: An economic approach. In *The economic dimensions of crime* (pp. 13–68). Springer.

Begus, K., Gliga, T., & Southgate, V. (2016). Infants' preferences for native speakers are associated with an expectation of information. *Proceedings of the National Academy of Sciences*, *113*(44), 12397–12402. doi: 10.1073/pnas.1603261113

Bekoff, M. (1972). The development of social interaction, play, and metacommunication in mammals: an ethological perspective. *The Quarterly Review of Biology*, *47*(4), 412–434. doi: 10.1086/407400

Bekoff, M. (1974). Social play and play-soliciting by infant canids. *American Zoologist*, *14*(1), 323–340. doi: 10.1093/icb/14.1.323

Bekoff, M. (1988). Motor training and physical fitness: Possible short-and long-term influences on the development of individual differences in behavior. *Developmental Psychobiology: The Journal of the International Society for Developmental Psychobiology*, *21*(6), 601–612. doi: 10.1002/dev.420210610

Bekoff, M., & Byers, J. A. (1985). The development of behavior from evolutionary and ecological perspectives in mammals and birds. In *Evolutionary biology* (pp. 215–286). Springer.

Bekoff, M., & Byers, J. A. (1998). *Animal play: Evolutionary, comparative and ecological perspectives*. Cambridge University Press.

Bennett, D., Bode, S., Brydevall, M., Warren, H., & Murawski, C. (2016). Intrinsic valuation of information in decision making under uncertainty. *PLoS Computational Biology*, *12*(7), e1005020. doi: 10.1371/journal.pcbi.1005020

Bennett-Pierre, G., & Gweon, H. (2018). Preschoolers consider expected task difficulty to decide what to do and whom to help. In *Proceedings of the 40th Annual Meeting of the Conference Cognitive Science Society.*

Berke, M., & Jara-Ettinger, J. (2021). *Thinking about thinking through inverse reasoning* (preprint). PsyArXiv. doi: 10.31234/osf.io/r25qn

Berlyne, D. (1960). *Conflict, arousal, and curiosity.* McGraw-Hill.

Berlyne, D. (1966). Curiosity and Exploration: Animals spend much of their time seeking stimuli whose significance raises problems for psychology. *Science*, *153*(3731), 25–33. doi: 10.1126/science.153.3731.25

Berlyne, D. (1969a). Arousal, reward and learning. *Annals of the New York Academy of Sciences*, *159*(3), 1059–1070. doi: 10.1111/j.1749-6632.1969.tb12997.x

Berlyne, D. (1969b). Laughter, humor, and play. In *Handbook of social psychology* (2nd ed., Vol. 3, pp. 795–852). Ma: Addison-Wesley.

Berman, C. (1982). Functions of play: first steps toward evolutionary explanation. *Behavioral and Brain Sciences*, *5*(1), 157–158. doi: 10.1017/s0140525x00010955

Bhui, R., Lai, L., & Gershman, S. J. (2021). Resource-rational decision making. *Current Opinion in Behavioral Sciences*, *41*, 15–21. doi: 10/gh74zd

Birch, S. A., Severson, R., & Baimel, A. (2020). Children's understanding of when a person's confidence and hesitancy is a cue to their credibility. *PLOS ONE*, *15*, 0227026. doi: 10/gr7zss

Bjorklund, D. F. (1997). The role of immaturity in human development. *Psychological Bulletin*, *122*(2), 153. doi: 10.1037/0033-2909.122.2.153

Blanchard, T. C., Hayden, B. Y., & Bromberg-Martin, E. S. (2015). Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron*, *85*(3), 602–614. doi: 10.1016/j.neuron.2014.12.050

Bloom, L., Merkin, S., & Wootten, J. (1982). " wh"-questions: Linguistic factors that contribute to the sequence of acquisition. *Child Development*, 1084–1092.

Bonawitz, E. B., & Lombrozo, T. (2012). Occam's rattle: Children's use of simplicity and probability to constrain inference. *Developmental Psychology*, *48*, 1156–1164. doi: 10.1037/a0026471

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

Bonawitz, E. B., van Schijndel, T. J., Friel, D., & Schulz, L. E. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, *64*(4), 215–234. doi: 10.1016/j.cogpsych.2011.12.002

Bornstein, M. H., Hahn, C.-S., & Suwalsky, J. T. (2013). Physically developed and exploratory young infants contribute to their own long-term academic achievement. *Psychological Science*, *24*(10), 1906–1917. doi: 10.1177/0956797613479974

Bramley, N. R., Schulz, E., Xu, F., & Tenenbaum, J. (2018). Learning as program induction. In C. Kalish, M. Rau, J. Zhu, & T. Rogers (Eds.), *Proceedings of the 40th annual conference of the cognitive science society* (pp. 5–6).

Bratman, M. (1987). *Intention, plans, and practical reason*. Harvard University Press, Cambridge.

Brennan, W. M., Ames, E. W., & Moore, R. W. (1966). Age differences in infants' attention to patterns of different complexities. *Science*, *151*(3708), 354–356. doi: 10.1126/science.151.3708.354

Breuggeman, J. (1978). The function of adult play in free-ranging macaca mulatta. In *Social play in primates* (pp. 169–191). Academic Press New York.

Bridgers, S., Buchsbaum, D., Seiver, E., Griffiths, T., & Gopnik, A. (2016). Children's causal inferences from conflicting testimony and observations. *Developmental Psychology*, *52*, 9–18. doi: 10/f75x8s

Bridgers, S., Jara-Ettinger, J., & Gweon, H. (2020). Young children consider the expected utility of others' learning to decide what to teach. *Nature Human Behaviour*, *4*(2), 144–152. doi: 10.1038/s41562-019-0748-6

Brockner, J. (1992). The Escalation of Commitment to a Failing Course of Action: Toward Theoretical Progress. *The Academy of Management Review*, *17*(1), 39–61. doi: 10/ff3t6f

Brockner, J., Rubin, J. Z., & Lang, E. (1981). Face-saving and entrapment. *Journal of Experimental Social Psychology*, *17*, 68–79. doi: 10/dn4bvp

Bromberg-Martin, E. S., & Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron*, *63*(1), 119–126. doi: 10.1016/j.neuron.2009.06.009

Bromberg-Martin, E. S., & Hikosaka, O. (2011). Lateral habenula neurons signal errors in the prediction of reward information. *Nature Neuroscience*, *14*(9), 1209–1216. doi: 10.1038/nn.2902

Brosseau-Liard, P., & Birch, S. (2010). I bet you know more and are nicer too!': What children infer from others' accuracy. *Developmental Science*, *13*, 772–778. doi: 10/dgwgn4

Brosseau-Liard, P., Cassels, T., & Birch, S. (2014). You seem certain but you were wrong before: Developmental change in preschoolers' relative trust in accurate versus confident speakers. *PLOS One*, *9:e108308*.

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

Buchsbaum, D., Bridgers, S., Skolnick Weisberg, D., & Gopnik, A. (2012). The power of possibility: causal learning, counterfactual reasoning, and pretend play. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *367*(1599), 2202–2212. doi: 10.1098/rstb.2012.0122

Burda, Y., Edwards, H., Pathak, D., Storkey, A., Darrell, T., & Efros, A. A. (2019). Large-scale study of curiosity-driven learning. In *ICLR.*

Burghardt, G. M. (2005). *The genesis of animal play: Testing the limits.* Mit Press.

Bustamante, A., Maldonado-Carreño, C., González, O. L., Navia, M. A., Restrepo, E. C., & Torres, D. F. (2020). "Who said we play to hit each other?" early childhood teachers' responses to children's transgressions. *Journal of Moral Education*, 1–17. doi: 10.1080/03057240.2019.1708710

Butler, L. P. (2020). The Empirical Child? A Framework for Investigating the Development of Scientific Habits of Mind. *Child Development Perspectives*, *14*(1), 34–40. doi: 10/gjr3qp

Butler, L. P., & Markman, E. M. (2012). Preschoolers use intentional and pedagogical cues to guide inductive inferences and exploration. *Child Development*, *83*(4), 1416–1428.

Butler, L. P., & Markman, E. M. (2014). Preschoolers use pedagogical cues to guide radical reorganization of category knowledge. *Cognition*, *130*(1), 116–127. doi: 10.1016/j.cognition.2013.10.002

Butler, L. P., Schmidt, M. F. H., Tavassolie, N. S., & Gibbs, H. M. (2018). Children's evaluation of verified and unverified claims. *Journal of Experimental Child Psychology*, *176*, 73–83. doi: 10.1016/j.jecp.2018.07.007

Byers, J. A. (1998). The biology of human play. *Child Development*, *69*(3), 599–600. doi: 10.1111/j.1467-8624.1998.tb06227.x

Byers, J. A., & Walker, C. (1995). Refining the Motor Training Hypothesis for the Evolution of Play. *The American Naturalist*, *146*(1), 25–40. doi: 10/fkg5hv

Cacioppo, J. T., & Petty, R. E. (1982). The need for cognition. *Journal of Personality and Social Psychology*, *42*(1), 116–131. doi: 10/ghc

Calhoun, A. J., Chalasani, S. H., & Sharpee, T. O. (2014). Maximally informative foraging by caenorhabditis elegans. *Elife*, *3*, e04220. doi: 10.7554/eLife.04220

Callanan, M. A., & Oakes, L. M. (1992). Preschoolers' questions and parents' explanations: Causal thinking in everyday activity. *Cognitive Development*, *7*(2), 213–233. doi: 10.1016/0885-2014(92)90012-g

Carey, S., & Spelke, E. (1994). Domain-specific knowledge and conceptual change. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the Mind* (pp. 169–200). Cambridge University Press. doi: 10.1017/cbo9780511752902.008

Caro, T. (1980). Effects of the mother, object play, and adult experience on predation in cats. *Behavioral and Neural Biology*, *29*(1), 29–51.

Caro, T. (1995). Short-term costs and correlates of play in cheetahs. *Animal Behavior*, *49*(2), 333–345. doi: 10.1006/anbe.1995.9999

Chan, C., & Tardif, T. (2013). Knowing better: The role of prior knowledge and culture in trust in testimony. *Developmental Psychology*, *49*(3), 591. doi: 10/f4p6q9

Charpentier, C. J., Bromberg-Martin, E. S., & Sharot, T. (2018). Valuation of knowledge and ignorance in mesolimbic reward circuitry. *Proceedings of the National Academy of Sciences*, *115*(31), E7255–e7264. doi: 10/gd2sfm

Chater, N., & Loewenstein, G. (2016). The under-appreciated drive for sense-making. *Journal of Economic Behavior & Organization*, *126*, 137–154. doi: 10/f8r6dj

Chater, N., & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Science*, *3*(2), 57–65. doi: 10.1016/s1364-6613(98)01273-x

Chen, E., Corriveau, K., & Harris, P. (2012). Children Trust a Consensus Composed of Outgroup Members-But Do Not Retain That Trust. *Child Development*, *84*(1), 269–282. doi: 10/f4nrmd

Cheng, S., Tang, N., An, W., Zhao, Y., Zhou, J., Shen, M., & Gao, T. (2022). *Intention beyond desire: Spontaneous intentional commitment regulates conflicting desires* (preprint). SSRN. doi: 10.2139/ssrn.4396194

Chi, M. T., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, *5*, 121–152. doi: 10/dv96v4

Chiang, L. S.-c. (2007). Daoist transcendence and tang literati identities in "records of mysterious anomalies" by niu sengru (780-848). *Chinese Literature: Essays, Articles, Reviews (CLEAR)*, *29*, 1–21.

Chitnis, R., Silver, T., Tenenbaum, J. B., Kaelbling, L. P., & Lozano-Pérez, T. (2021). GLIB: Efficient Exploration for Relational Model-Based Reinforcement Learning via Goal-Literal Babbling. In *The 35th AAAI Conference on Artificial Intelligence* (pp. 11782–11791).

Chouinard, M. M. (2007). Children's questions: A mechanism for cognitive development. *Monogr Soc Res Child Dev*, *72*(1), 1–112.

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

Chu, J., & Schulz, L. E. (2020b). Play, Curiosity, and Cognition. *Annual Review of Developmental Psychology*, *2*(1), 317–343. doi: 10.1146/annurev-devpsych-070120-014806

Chu, J., & Schulz, L. E. (2021). Children selectively endorse speculative conjectures. *Child Development*, *92*(6), e1342–1360. doi: 10.1111/cdev.13647

Chu, J., & Schulz, L. E. (in press). Not Playing by the Rules: Exploratory Play, Rational Action, and Efficient Search. *Open Mind*. doi: https://doi.org/10.31234/osf.io/9yra2

Civan, A., Teller, D. Y., & Palmer, J. (2005). Relations among spontaneous preferences, familiarized preferences, and novelty effects: Measurements with forced-choice techniques. *Infancy*, *7*(2), 111–142. doi: 10.1207/s15327078in0702_1

Clark, E. V. (1990). On the pragmatics of contrast. *Journal of child language*, *17*(2), 417–431. doi: 10.1017/s0305000900013842

Clément, F., Koenig, M. A., & Harris, P. (2004). The ontogenesis of trust. *Mind & Language*, *19*(4), 360–379. doi: 10/bw7t4p

Cohen, L. B. (1972). Attention-getting and attention-holding processes of infant visual preferences. *Child Development*, *43*(3), 869–879. doi: 10/fjv4rj

Colas, C., Karch, T., Sigaud, O., & Oudeyer, P.-Y. (2022). Autotelic Agents with

Intrinsically Motivated Goal-Conditioned Reinforcement Learning: a Short Survey. *Journal of Artificial Intelligence Research*, *74*.

Colliver, Y., & Fleer, M. (2016). ‘I already know what I learned’: young children’s perspectives on learning through play. *Early Child Development and Care*, *186*(10), 1559–1570. doi: 10.1080/03004430.2015.1111880

Connolly, K., & Dalgleish, M. (1989). The emergence of a tool-using skill in infancy. *Developmental Psychology*, *25*(6), 894. doi: 10.1037/0012-1649.25.6.894

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

Cooper, R. M. (1974). The control of eye fixation by the meaning of spoken language: a new methodology for the real-time investigation of speech perception, memory, and language processing. *Cognitive psychology*.

Cordoni, G. (2009). Social play in captive wolves (canis lupus): not only an immature affair. *Behaviour*, *146*(10), 1363–1385. doi: 10.1163/156853909x427722

Correia, S. P., Dickinson, A., & Clayton, N. S. (2007). Western scrub-jays anticipate future needs independently of their current motivational state. *Current Biology*, *17*(10), 856–861. doi: 10/d3fmn8

Corriveau, K., & Harris, P. (2010). Preschoolers (sometimes) defer to the majority in making simple perceptual judgments. *Developmental Psychology*, *46*(2), 437–445. doi: 10/dmx7r6

Corriveau, K., & Kurkul, K. (2014). *Why Does Rain Fall?": Children prefer to learn from an informant who uses noncircular explanations.*

Csibra, G. (2016). Statistical treatment of looking-time data. *Developmental Psychology*, *52*(4), 521. doi: 10/f8hghz

Csibra, G., Bíró, S., Koós, O., & Gergely, G. (2003). One-year-old infants use teleological representations of actions productively. *Cognitive Science*, *27*(1), 111–133. doi: 10.1207/s15516709cog2701_4

Csikszentmihalyi, M., & Csikzentmihaly, M. (1990). *Flow: The psychology of optimal experience* (Vol. 1990). Harper & Row New York.

Daddaoua, N., Lopes, M., & Gottlieb, J. (2016). Intrinsically motivated oculomotor exploration guided by uncertainty reduction and conditioned reinforcement in non-human primates. *Scientific Reports*, *6*(1), 20202. doi: 10.1038/srep20202

Dall, S., Giraldeau, L., Olsson, O., Mcnamara, J., & Stephens, D. (2005). Information and its use by animals in evolutionary ecology. *Trends in Ecology & Evolution*, *20*(4), 187–193. doi: 10.1016/j.tree.2005.01.010

Dan, O., Leshkowitz, M., & Hassin, R. R. (2020). On clickbaits and evolution: curiosity from urge and interest. *Current Opinion in Behavioral Sciences*, *35*, 150–156. doi: 10/grwmvq

Danckert, J., & Elpidorou, A. (2023). In search of boredom: beyond a functional account. *Trends in Cognitive Sciences*.

Danovitch, J., & Keil, F. (2004). Should you ask a fisherman or a biologist?: Developmental shifts in ways of clustering knowledge. *Child Development*, *75*(3), 918–931. doi: 10/bz52zk

Daw, N. D., O’doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical

substrates for exploratory decisions in humans. *Nature*, *441*(7095), 876–879. doi: 10.1038/nature04766
Dawkins, M. S. (2006). A user's guide to animal welfare science. *Trends in Ecology & Evolution*, *21*(2), 77–82. doi: 10.1016/j.tree.2005.10.017
Day, S. X., & Rounds, J. (1998). Universality of vocational interest structure among racial and ethnic minorities. *American Psychologist*, *53*(7), 728–736. doi: 10/fd5fg5
Dayan, P., & Niv, Y. (2008). Reinforcement learning: The Good, The Bad and The Ugly. *Current Opinion in Neurobiology*, *18*(2), 185–196. doi: 10.1016/j.conb.2008.08.003
Dayan, P., & Sejnowski, T. J. (1996). Exploration bonuses and dual control. *Machine Learning*, *25*(1), 5–22.
Deakos, M. H., Branstetter, B. K., Mazzuca, L., Fertl, D., & Mobley Jr, J. R. (2010). Two unusual interactions between a bottlenose dolphin (tursiops truncatus) and a humpback whale (megaptera novaeangliae) in hawaiian waters. *Aquatic Mammals*, *36*(2). doi: 10.1578/am.36.2.2010.121
de Almeida Soares, D., von Hofsten, C., & Tudella, E. (2012). Development of exploratory behavior in late preterm infants. *Infant Behavior and Development*, *35*(4), 912–915. doi: 10/f22qb2
de Campos, A. C., da Costa, C. S. N., Savelsbergh, G. J., & Rocha, N. A. C. F. (2013). Infants with down syndrome and their interactions with objects: development of exploratory actions after reaching onset. *Research in Developmental Disabilities*, *34*(6), 1906–1916. doi: 10.1016/j.ridd.2013.03.001
Deci, E. L., & Ryan, R. M. (1980). The empirical exploration of intrinsic motivational processes. In *Advances in experimental social psychology* (Vol. 13, pp. 39–80). Elsevier.
Delfour, F., & Aulagnier, S. (1997). Bubbleblow in beluga whales (delphinapterus leucas): a play activity? *Behavioural Processes*, *40*(2), 183–186. doi: 10/d5nb9j
Dember, W. N., & Earl, R. W. (1957). Analysis of exploratory, manipulatory, and curiosity behaviors. *Psychological Review*, *64*(2), 91–96. doi: 10.1037/h0046861
Dennett, D. C. (1984). Cognitive wheels: The frame problem of ai. *Minds, machines and evolution*, 129–151.
Devinney, B., Berman, C., & Rasmussen, K. (2003). Individual differences in response to sibling birth among free-ranging yearling rhesus monkeys (macaca mulatta) on cayo santiago. *Behaviour*, *140*(7), 899–924. doi: 10.1163/156853903770238373
De Waal, F. B. (1986). The integration of dominance and social bonding in primates. *The Quarterly Review of Biology*, *61*(4), 459–479. doi: 10.1086/415144
Dewey, J. (1916). What pragmatism means by practical. In *Essays in experimental logic* (pp. 303–329).
Diamond, A. (2013). Executive functions. *Annual review of psychology*, *64*, 135–168.
Dietz, G., Landay, J. A., & Gweon, H. (2019). Building blocks of computational thinking: Young children's developing capacities for problem decomposition. In *Cogsci* (pp. 1647–1653).
DiYanni, C., Corriveau, K., Kurkul, K., Nasrini, J., & Nini, D. (2015). The role of consensus and culture in children's imitation of inefficient actions. *Journal of*

Experimental Child Psychology, *137*, 99–110. doi: 10/cfxb

Doebel, S., Barker, J. E., Chevalier, N., Michaelson, L. E., Fisher, A. V., & Munakata, Y. (2017). Getting ready to use control: Advances in the measurement of young children's use of proactive control. *PloS one*, *12*(4), e0175072.

Doebel, S., & Lillard, A. S. (2023). How does play foster development? a new executive function perspective. *Developmental Review*, *67*, 101064. doi: 10.31234/osf.io/q5mdc

Dolhinow, P. (1999). Play: a critical process in the developmental system. *The nonhuman primates*, 231–236.

Dominy, N. J., Garber, P. A., Bicca-Marques, J. C., & Azevedo-Lopes, M. A. d. O. (2003). Do female tamarins use visual cues to detect fruit rewards more successfully than do males? *Animal Behaviour*, *66*(5), 829–837. doi: 10.1006/anbe.2003.2288

Drea, C. M., Hawk, J. E., & Glickman, S. E. (1996). Aggression decreases as play emerges in infant spotted hyaenas: preparation for joining the clan. *Animal Behaviour*, *51*(6), 1323–1336. doi: 10.1006/anbe.1996.0136

Dubey, R., Griffiths, T. L., & Lombrozo, T. (2022). If it's important, then I'm curious: Increasing perceived usefulness stimulates curiosity. *Cognition*, *226*, 105193. doi: 10/grwms8

Dunbar, R. I. (2012). On the evolutionary function of song and dance. *Music, language, and human evolution*, 201–14. doi: 10.1093/acprof:osobl/9780199227341.003.0008

Edwards, C. P. (2000). Children's play in cross-cultural perspective: A new look at the six cultures study. *Cross-Cultural Research*, *34*(4), 318–338. doi: 10.1177/106939710003400402

Einav, S., & Robinson, E. (2011). When being right is not enough: Four-year-olds distinguish knowledgeable informants from merely accurate informants. *Psychological Science*, *22*, 1250–1253. doi: 10/d57rk9

Eisenberger, R. (1992). Learned industriousness. *Psychological Review*, *99*(2), 248–267. doi: 10/frcw7r

Elashi, F., & Mills, C. (2014). Do children trust based on group membership or prior accuracy? The role of novel group membership in children's trust decisions. *Journal of Experimental Child Psychology*, *128*, 88–104. doi: 10/f6kzms

Ellis, K., Wong, C., Nye, M., Sablé-Meyer, M., Morales, L., Hewitt, L., … Tenenbaum, J. B. (2021). DreamCoder: bootstrapping inductive program synthesis with wake-sleep library learning. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation* (pp. 835–850). Acm. doi: 10.1145/3453483.3454080

Emlen, S. T. (1975). Migration: orientation and navigation. In *Avian Biology* (pp. 129–219). Elsevier. doi: 10.1016/b978-0-12-249405-5.50011-2

Ervin-Tripp, S. (1970). Discourse agreement: How children answer questions. In J. Hayes (Ed.), *Cognition and the development of knowledge.* Wiley.

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

Fagan, R. M. (1981). *Animal play behavior.* Ny: Oxford University Press.

Fagone, J. (2017). *The Woman Who Smashed Codes: A True Story of Love, Spies, and the Unlikely Heroine Who Outwitted America's Enemies*. New York: Dey Street Books.

Fagot, B., & Kavanagh, K. (1991). Play as a diagnostic tool with physically abusive parents and their children. *Play Diagnosis and Assessment. New York: Wiley*, 203–18.

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

Ferrigno, S., Cheyette, S. J., Piantadosi, S. T., & Cantlon, J. F. (2020). Recursive sequence generation in monkeys, children, U.S. adults, and native Amazonians. *Science Advances*, *6*(26), eaaz1002. doi: 10.1126/sciadv.aaz1002

Fireman, G., Kose, G., & Solomon, M. J. (2003). Self-observation and learning: The effect of watching oneself on problem solving performance. *Cognitive Development*, *18*(3), 339–354. doi: 10.1016/s0885-2014(03)00038-8

Fiser, J., & Aslin, R. N. (2002). Nonlinear partial differential equations and applications: From the cover: Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences*, *99*(24), 15822–15826. doi: 10.1073/pnas.232472899

Fisher, K. R., Hirsh-Pasek, K., Newcombe, N., & Golinkoff, R. M. (2013). Taking shape: Supporting preschoolers' acquisition of geometric knowledge through guided play. *Child Development*, *84*(6), 1872–1878. doi: 10.1111/cdev.12091

FitzGibbon, L., Lau, J. K. L., & Murayama, K. (2020). The seductive lure of curiosity: information as a motivationally salient reward. *Current Opinion in Behavioral Sciences*, *35*, 21–27. doi: 10/gqn2zg

Flavell, J. (1988). The development of children's knowledge about the mind: From cognitive connections to mental representations. In J. Astington, P. Harris, & D. Olson (Eds.), *Developing theories of mind.* Cambridge University Press.

Florensa, C., Held, D., Geng, X., & Abbeel, P. (2017). Automatic goal generation for reinforcement learning agents. *Icml*.

Florensa, C., Held, D., Geng, X., & Abbeel, P. (2018). Automatic Goal Generation for Reinforcement Learning Agents. In J. G. Dy & A. Krause (Eds.), *Proceedings of the 35th International Conference on Machine Learning* (Vol. 80, pp. 1514–1523). Pmlr.

Forbus, K. D., Liang, C., & Rabkina, I. (2017). Representation and Computation in Cognitive Models. *Topics in Cognitive Science*, *9*(3), 694–718. doi: 10.1111/tops.12277

Forestier, S., Mollard, Y., & Oudeyer, P.-Y. (2017). Intrinsically Motivated Goal Exploration Processes with Automatic Curriculum Learning. *arXiv*, *1708.02190*.

Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, *336*(6084), 998–998.

Fraser, D., & Duncan, I. J. (1998). Pleasures, pains and animal welfare: toward a natural history of affect. *Animal Welfare*, *7*, 383–396.

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. doi: 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. doi: 10.1080/15248372.2015.1098649

Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, *6*(4), 187–214. doi: 10/gf393f

Fusaro, M., Corriveau, K., & Harris, P. (2011). The good, the strong, and the accurate: Preschoolers' evaluations of informant attributes. *Journal of Experimental Child Psychology*, *110*, 561–574. doi: 10/frqb87

Gaines, M. S., & McClenaghan Jr, L. R. (1980). Dispersal in small mammals. *Annual Review of Ecology and Systematics*, *11*(1), 163–196. doi: 10/dnfmbs

Galyer, K. T., & Evans, I. M. (2001). Pretend Play and the Development of Emotion Regulation in Preschool Children. *Early Child Development and Care*, *166*(1), 93–108. doi: 10/fsd4cr

Garland, H. (1990). Throwing good money after bad: The effect of sunk costs on the decision to escalate commitment to an ongoing project. *Journal of Applied Psychology*, *75*(6), 728.

Gendler, T. S. (2000). The Puzzle of Imaginative Resistance. *The Journal of Philosophy*, *97*(2), 55–81. doi: 10.2307/2678446

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: the naive theory of rational action. *Trends Cognition Sciences*, *7*(7), 287–292. doi: 10.1016/s1364-6613(03)00128-1

Gergely, G., Nádasdy, Z., Csibra, G., & Bíró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, *56*(2), 165–193. doi: 10.1016/0010-0277(95) 00661-h

Gerken, L., Balcomb, F. K., & Minton, J. L. (2011). Infants avoid 'labouring in vain'by attending more to learnable than unlearnable linguistic patterns. *Developmental Science*, *14*(5), 972–979. doi: 10.1111/j.1467-7687.2011.01046.x

Gershman, S. J., & Daw, N. D. (2017). Reinforcement Learning and Episodic Memory in Humans and Animals: An Integrative Framework. *Annual Review of Psychology*, *68*(1), 101–128. doi: 10/gdz7n8

Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, *349*(6245), 273–278. doi: 10.1126/science.aac6076

Gigerenzer, G. (2008). *Rationality for mortals: How people cope with uncertainty*. Oxford University Press.

Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in cognitive science*, *1*(1), 107–143.

Gilpin, A. T., Brown, M. M., & Pierucci, J. M. (2015). Relations Between Fantasy Orientation and Emotion Regulation in Preschool. *Early Education and Development*, *26*(7), 920–932. doi: 10/gf7nz5

Glymour, C. N. (2015). *Thinking things through: an introduction to philosophical issues and achievements* (Second edition. ed.). Cambridge, Massachusetts: MIT Press.

Goddu, M., Rule, J., Bonawitz, E., Gopnik, A., & Ullman, T. (2022). *Fun isn't easy: Children optimize for difficulty when "playing for fun" vs."playing to win".*

Golinkoff, D. G. S. R. M., Hirsh-Pasek, K., et al. (2006). *Play= learning: How play motivates and enhances children's cognitive and social-emotional growth.* Oxford University Press.

Goodhall, N., & Atkinson, C. (2019). How do children distinguish between 'play' and 'work'? Conclusions from the literature. *Early Child Development and Care*, *189*(10), 1695–1708. doi: 10.1080/03004430.2017.1406484

Goodman, N. (1955/1983). *Fact, fiction, and forecast.* Bobbs-Merrill.

Goodman, N. A., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A rational analysis of rule-based concept learning. *Cognitive Science*, *32*(1), 108–154. doi: 10/bssmgd

Gopnik, A., & Walker, C. M. (2013). Considering counterfactuals the relationship between causal learning and pretend play. *American Journal of Play*, *6*(1), 15–28.

Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, bayesian learning mechanisms, and the theory theory. *Psychological Bulletin*, *138*(6), 1085–1108.

Gordon, G. (2020). Infant-inspired intrinsically motivated curious robots. *Current Opinion in Behavioral Sciences*, *35*, 28–34. doi: 10.1016/j.cobeha.2020.05.010

Gosso, Y. (2010). Play in different cultures. In P. Smith (Ed.), *Children at play* (pp. 80–98). West Sussex, UK: Wiley Blackwell.

Gosso, Y., e Morais, M. d. L. S., & Otta, E. (2007). Pretend play of brazilian children: A window into different cultural worlds. *Journal of Cross-Cultural Psychology*, *38*(5), 539–558. doi: 10.1177/0022022107305237

Gottlieb, J., & Oudeyer, P.-Y. (2018). Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience*, *19*(12), 758–770. doi: 10.1038/s41583-018-0078-0

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

Gould, S. J., & Lewontin, R. C. (1979). The spandrels of san marco and the panglossian paradigm: a critique of the adaptationist programme. *Proceedings of the Royal Society of London B: Biological Sciences*, *205*(1161), 581–598. doi: 10.1098/rspb.1979.0086

Goupil, L., & Kouider, S. (2016). Behavioral and Neural Indices of Metacognitive Sensitivity in Preverbal Infants. *Current biology: CB*, *26*(22), 3038–3045. doi: 10/f9c3xt

Goupil, L., & Kouider, S. (2019). Developing a Reflective Mind: From Core Metacognition to Explicit Self-Reflection. *Current Directions in Psychological Science*, *28*(4), 403–408. doi: 10/gm2c6d

Gray, W. D., & Banerjee, S. (2021). Constructing expertise: Surmounting performance plateaus by tasks, by tools, and by techniques. *Topics in Cognitive Science*, *13*(4), 610–665.

Groos, K. (1898). *The play of animals.* NY: Appleton.

Groos, K. (1901). *The play of man.* NY: Appleton.
Gulick, L. H. (1920). *A Philosophy of Play.* McGrath Publishing Company.
Gureckis, T. M., & Markant, D. B. (2012). Self-directed learning: A cognitive and computational perspective. *Perspectives on Psychological Science*, *7*(5), 464–481. doi: 10.1177/1745691612454304
Gusinde, M. (1931). *Die feuerland indianer: Ergebnisse meiner vier forschungsreisen. [the indians of the land of fire. results from my four research journeys.].* Verlag der internationalen Zeitschrift Anthropos.
Gweon, H., Pelton, H., Konopka, J. A., & Schulz, L. E. (2014). Sins of omission: Children selectively explore when teachers are under-informative. *Cognition*, *132*(3), 335–341. doi: 10.1016/j.cognition.2014.04.013
Gweon, H., & Schulz, L. E. (2011). 16-month-olds rationally infer causes of failed actions. *Science*, *332*(6037), 1524–1524. doi: 10.1126/science.1204493
Gweon, H., & Schulz, L. E. (2018). From exploration to instruction: Children learn from exploration and tailor their demonstrations to observers' goals and competence. *Child Development*, *90*(1), e148–e164. doi: 10.1111/cdev.13059
Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences*, *107*(20), 9066–9071. doi: 10.1073/pnas.1003095107
Haber, N., Mrowca, D., Wang, S., Li, F.-F., & Yamins, D. L. K. (2018). Learning to Play with Intrinsically-Motivated, Self-Aware Agents. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems* (pp. 8398–8409). Red Hook, NY, USA: Curran Associates Inc.
Haith, M. M. (1980). *Rules that babies look by: The organization of newborn visual activity.* Lawrence Erlbaum Associates.
Hall, K. (1968). Behaviour and ecology of the wild patas monkey, erythrocebus patas. In *Primates: studies in adaptation and variability.* (pp. 32–119). Ny: Holt, Rinehart, and Winston.
Hall, S. L., & Bradshaw, J. W. (1998). The influence of hunger on object play by adult domestic cats. *Applied Animal Behaviour Science*, *58*(1-2), 143–150. doi: 10.1016/s0168-1591(97)00136-6
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. doi: 10.1111/j.1467-7687.2010.00951.x
Harcourt, R. (1991). Survivorship costs of play in the south american fur seal. *Animal behaviour*. doi: 10.1016/s0003-3472(05)80055-7
Harris, P. L. (2000). *Understanding children's worlds: The work of the imagination.* Blackwell.
Harris, P. L. (2021). Early Constraints on the Imagination: The Realism of Young Children. *Child Development*, *92*(2), 466–483. doi: 10.1111/cdev.13487
Harris, P. L., & Kavanaugh, R. D. (1993). Young Children's Understanding of Pretense. *Monographs of the Society for Research in Child Development*, *58*(1), i–107. doi: 10.2307/1166074
Harris, P. L., Koenig, M. A., Corriveau, K. H., & Jaswal, V. K. (2018). Cognitive foundations of learning from testimony. *Annual Review of Psychology*, *69*,

251–273. doi: 10/gdrbzv
Harvey, P., & Victoravich, L. M. (2009). The influence of forward-looking antecedents, uncertainty, and anticipatory emotions on project escalation. *Decision Sciences*, *40*(4), 759–782.
Hatano, A., Ogulmus, C., Shigemasu, H., & Murayama, K. (2022). Thinking about thinking: People underestimate how enjoyable and engaging just waiting is. *Journal of Experimental Psychology: General*, *151*(12), 3213–3229. doi: 10/grwmwb
Hausberger, M., Fureix, C., Bourjade, M., Wessel-Robert, S., & Richard-Yris, M.-A. (2012). On the significance of adult play: what does social play tell us about adult horse welfare? *Naturwissenschaften*, *99*(4), 291–302. doi: 10.1007/s00114-012-0902-8
Hauser, M. D., Chomsky, N., & Fitch, W. T. S. (2002). The faculty of language: what is it, who has it, and how did it evolve? *Science*, *298*(5598), 1569–79. doi: 10.1126/science.298.5598.1569
Healey, C. J. (1990). *Maring hunters and traders: production and exchange in the papua new guinea highlands* (Vol. 8). Univ of California Press.
Held, S. D., & Špinka, M. (2011). Animal play and animal welfare. *Animal Behaviour*, *81*(5), 891–899. doi: 10/b53nd2
Henrich, J. P. (2016). *The secret of our success: how culture is driving human evolution, domesticating our species, and making us smarter*. Princeton: Princeton university press.
Hilbert, M. (2012). Toward a synthesis of cognitive biases: how noisy information processing can bias human decision making. *Psychological bulletin*, *138*(2), 211. doi: 10/d6tnkd
Hoehl, S., Keupp, S., Schleihauf, H., McGuigan, N., Buttelmann, D., & Whiten, A. (2019). 'Over-imitation': A review and appraisal of a decade of research. *Developmental Review*, *51*, 90–108. doi: 10/gfspt3
Holdaway, C., Bear, D. M., Radwan, S. F., Frank, M. C., Yamins, D. L., & Fan, J. E. (2021). Measuring and predicting variation in the interestingness of physical structures. In *Proceedings of the 43rd annual meeting of the cognitive science society.*
Holmes, W. G. (1995). The ontogeny of littermate preferences in juvenile golden-mantled ground squirrels: effects of rearing and relatedness. *Animal Behaviour*, *50*(2), 309–322. doi: 10.1006/anbe.1995.0247
Holton, R. (2004). Rational resolve. *The Philosophical Review*, *113*(4), 507–535.
Horner, V., & Whiten, A. (2005). Causal knowledge and imitation/emulation switching in chimpanzees (Pan troglodytes) and children (Homo sapiens). *Animal cognition*, *8*(3), 164–181. doi: 10/d6r3hn
Houthooft, R., Chen, X., Chen, X., Duan, Y., Schulman, J., De Turck, F., & Abbeel, P. (2016). VIME: Variational Information Maximizing Exploration. In *Advances in Neural Information Processing Systems* (Vol. 29, pp. 1109–1117). Curran Associates, Inc.
Hsee, C. K., & Ruan, B. (2016). The Pandora effect: The power and peril of curiosity. *Psychological science*, *27*(5), 659–666. doi: 10/f8n88z

Hughes, C., & Dunn, J. (1997). “Pretend you didn’t know”: Preschoolers’ talk about mental states in pretend play. *Cognitive development*, *12*(4), 477–497. doi: 10/b3s6db

Huh, M., Grossmann, I., & Friedman, O. (2019). Children show reduced trust in confident advisors who are partially informed. *Cognitive Development*, *50*, 49–55. doi: 10/gr7zst

Humle, T. (2006). Ant dipping in chimpanzees: An example of how microecological variables, tool use, and culture reflect the cognitive abilities of chimpanzees. In *Cognitive development in chimpanzees* (pp. 452–475). Springer.

Hunt, L. T., Daw, N. D., Kaanders, P., MacIver, M. A., Mugan, U., Procyk, E., ... Kolling, N. (2021). Formalizing planning and information search in naturalistic decision-making. *Nature Neuroscience*, *24*(8), 1051–1064. doi: 10.1038/s41593-021-00866-w

Ingebrigtsen, A. (1929). Whales caught in the north atlantic and other seas. *Conseil Permanent International pour l’Exploration de la Mer. Rapports et Proces-Verbaux des Reunions*, *56*, 123–135.

Inoue-Nakamura, N., & Matsuzawa, T. (1997). Development of stone tool use by wild chimpanzees (pan troglodytes). *Journal of Comparative Psychology*, *111*(2), 159. doi: 10/fjp3xg

Inzlicht, M., Shenhav, A., & Olivola, C. Y. (2018). The Effort Paradox: Effort Is Both Costly and Valued. *Trends in Cognitive Sciences*, *22*(4), 337–349. doi: 10/gc9qnq

Jacobs, C., Lopez-Brau, M., & Jara-Ettinger, J. (2021). What happened here? children integrate physical reasoning to infer actions from indirect evidence. *Proceedings of the 43rd Annual Conference of the Cognitive Science Society*.

James, W. (1890). *The Principles of Psychology* (Vol. 1).

James, W. (1907). *Pragmatism*. Project Gutenberg.

Janacsek, K., Fiser, J., & Nemeth, D. (2012). The best time to acquire new skills: Age-related differences in implicit sequence learning across the human lifespan. *Developmental science*, *15*(4), 496–505. doi: 10.1111/j.1467-7687.2012.01150.x

Jara-Ettinger, J., Floyd, S., Tenenbaum, J. B., & Schulz, L. E. (2017). Children understand that agents maximize expected utilities. *Journal of Experimental Psychology: General*, *146*(11), 1574–1585. doi: 10.1037/xge0000345

Jara-Ettinger, J., Gibson, E., Kidd, C., & Piantadosi, S. (2016). Native Amazonian children forego egalitarianism in merit-based tasks when they learn to count. *Developmental Science*, *19*(6), 1104–1110. doi: 10.1111/desc.12351

Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The Naïve Utility Calculus: Computational Principles Underlying Commonsense Psychology. *Trends in Cognitive Sciences*, *20*(10), 785. doi: 10.1016/j.tics.2016.08.007

Jara-Ettinger, J., Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2015). Children’s understanding of the costs and rewards underlying rational action. *Cognition*, *140*, 14–23. doi: 10.1016/j.cognition.2015.03.006

Jaswal, V. K. (2004). Don’t believe everything you hear: Preschoolers’ sensitivity to speaker intent in category induction. *Child Development*, *7*, 1871–1885. doi:

10/cstr85
Jaswal, V. K., & Malone, L. S. (2007). Turning believers into skeptics: 3-year-olds' sensitivity to cues to speaker credibility. *Journal of Cognition and Development*, *8*(3), 263–283. doi: 10.1080/15248370701446392
Jay, P. (1963). Mother-infant relationships in langurs. In H. Rheingold (Ed.), *Maternal behavior in mammals* (pp. 282–304). Ny: John Wiley & Sons.
Jean, A., Daubert, E. N., Yu, Y., Shafto, P., & Bonawitz, E. (2019). Pedagogical questions empower exploration. In *Proceedings of the 41st annual conference of the cognitive science society* (pp. 485–491).
Jirout, J., & Klahr, D. (2012). Children's scientific curiosity: In search of an operational definition of an elusive concept. *Developmental Review*, *32*(2), 125–160.
Johnston, A. M., Johnson, S. G., Koven, M. L., & Keil, F. C. (2016). Little bayesians or little einsteins? probability and explanatory virtue in children's inferences. *Developmental Science*, *20*(6), e12483. doi: 10.1111/desc.12483
Jones, S. S., Collins, K., & Hong, H.-W. (1991). An Audience Effect on Smile Production in 10-Month-Old Infants. *Psychological Science*, *2*(1), 45–49. doi: 10.1111/j.1467-9280.1991.tb00095.x
Jurasz, C. M., & Jurasz, V. P. (1979). Feeding modes of the humpback whale, megaptera novaeangliae, in southeast alaska. *Scientific Reports of the Whales Research Institute*, *31*, 69–83.
Kabadayi, C., & Osvath, M. (2017). Ravens parallel great apes in flexible planning for tool-use and bartering. *Science*, *357*(6347), 202–204. doi: 10/gbptwv
Kaelbling, L. P. (1993). Learning to Achieve Goals. In *Proceedings of IJCAI-93* (pp. 1094–1098). Morgan Kaufmann.
Kahneman, D., & Tversky, A. (1982). The Psychology of Preferences. *Scientific American*, *246*(1), 160–173. doi: 10/fqj9th
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. doi: 10.1111/cdev.12086
Kaur, M., Srinivasan, S. M., & Bhat, A. N. (2015). Atypical object exploration in infants at-risk for autism during the first year of lifer. *Frontiers in Psychology*, *6*. doi: 10/gq6jbm
Kavanaugh, R., & Harris, P. L. (1999). Pretense and counterfactual thought in young children. In L. Balter & C. S. Tamis-LeMonda (Eds.), *Child psychology: A handbook of contemporary issues* (pp. 158–176). Psychology Press.
Kavšek, M. (2004). Infant perception of object unity in static displays. *International Journal of Behavioral Development*, *28*(6), 538–545. doi: 10.1080/01650250444000252
Kavšek, M., & Bornstein, M. H. (2010). Visual habituation and dishabituation in preterm infants: A review and meta-analysis. *Research in Developmental Disabilities*, *31*(5), 951–975. doi: 10.1016/j.ridd.2010.04.016
Kenward, B., Karlsson, M., & Persson, J. (2011). Over-imitation is better explained by norm learning than by distorted causal learning. *Proc. R. Soc. B.*, *278*, 1239–1246.
Keupp, S., Behne, T., & Rakoczy, H. (2013). Why do children overimitate? Normativity

is crucial. *Journal of Experimental Child Psychology*, *116*(2), 392–406. doi: 10/f5fht9
Keupp, S., Behne, T., & Rakoczy, H. (2018). The Rationality of (Over)imitation. *Perspectives on Psychological Science*, *13*(6), 678–687. doi: 10/gq6jcp
Kidd, C., & Hayden, B. Y. (2015). The Psychology and Neuroscience of Curiosity. *Neuron*, *88*(3), 449–460. doi: 10.1016/j.neuron.2015.09.010
Kidd, C., Palmeri, H., & Aslin, R. N. (2013). Rational snacking: Young children's decision-making on the marshmallow task is moderated by beliefs about environmental reliability. *Cognition*, *126*(1), 109–114. doi: 10.1016/j.cognition.2012.08.004
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. doi: 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. doi: 10.1111/cdev.12263
Kim, E. Y., & Song, H.-j. (2015). Six-month-olds actively predict others' goal-directed actions. *Cognitive Development*, *33*, 1–13. doi: 10.1016/j.cogdev.2014.09.003
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. doi: 10.1037/0096-1523.18.2.385
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. doi: 10.1111/1467-7687.00080
Kinney, D. K., & Kagan, J. (1976). Infant attention to auditory discrepancy. *Child Development*, 155–164. doi: 10.2307/1128294
Kinzler, K., Corriveau, K., & Harris, P. (2011). Children's selective trust in native-accented speakers. *Developmental Science*, *14*, 106–111. doi: 10/dbv56d
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. doi: 10.1016/s0010-0277(02)00004-5
Klinger, E. (1969). Development of imaginative behavior: Implications of play for a theory of fantasy. *Psychological Bulletin*, *72*(4), 277–298. doi: 10/cvjdv5
Klyubin, A. S., Polani, D., & Nehaniv, C. L. (2005). Empowerment: a universal agent-centric measure of control. In *2005 ieee congress on evolutionary computation* (Vol. 1, pp. 128–135 Vol.1). doi: 10.1109/cec.2005.1554676
Koenig, M. A., Clement, F., & Harris, P. L. (2004). Trust in testimony: children's use of true and false statements. *Psychological Science*, *14*(10), 694–698. doi: 10.1111/j.0956-7976.2004.00742.x
Koenig, M. A., & Echols, C. (2003). Infants' understanding of false labeling events: The referential roles of words and the speakers who use them. *Cognition*, *87*, 179–208. doi: 10/c3wp7c
Koenig, M. A., & Harris, P. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, *76*, 1261–1277. doi: 10/dw4t27
Koenig, M. A., & Jaswal, V. K. (2011). Characterizing children's expectations about

expertise and incompetence: Halo or pitchfork effects? *Child Development*, *82*, 1634–1647. doi: 10/dvncz2

Koenig, M. A., Tiberius, V., & Hamlin, J. K. (2019). Children's Judgments of Epistemic and Moral Agents: From Situations to Intentions. *Perspectives on Psychological Science*. doi: 10.1177/1745691618805452

Koenig, M. A., & Woodward, A. (2010). Sensitivity of 24-month-olds to the prior inaccuracy of the source: possible mechanisms. *Developmental Psychology*, *46*, 815–826. doi: 10/ddm9kj

Kominsky, J. F., Langthorne, P., & Keil, F. C. (2016). The better part of not knowing: Virtuous ignorance. *Developmental Psychology*, *52*(1), 31–45. doi: 10.1037/dev0000065

Kopp, C. B., & Vaughn, B. E. (1982). Sustained attention during exploratory manipulation as a predictor of cognitive competence in preterm infants. *Child Development*, 174–182. doi: 10.2307/1129650

Koterba, E. A., Leezenbaum, N. B., & Iverson, J. M. (2014). Object exploration at 6 and 9 months in infants with and without risk for autism. *Autism*, *18*(2), 97–105. doi: 10.1177/1362361312464826

Kuczaj, S. A., & Eskelinen, H. C. (2014). Why do dolphins play. *Animal Behavior and Cognition*, *1*(2), 113–127. doi: 10/gq6jbc

Kurkul, K., & Corriveau, K. (2017). Question, Explanation, Follow-Up: A mechanism for learning from others? *Child Development*, *89*, 280–294. doi: 10/gcvwf9

Kushnir, T., & Gopnik, A. (2005). Young children infer causal strength from probabilities and interventions. *Psychological science*, *16*(9), 678–683. doi: 10/b73cnt

Kushnir, T., Vredenburgh, C., & Schneider, L. (2013). Who can help me fix this toy?" The distinction between causal knowledge and word knowledge guides preschoolers' selective requests for information. *Developmental Psychology*, *49*, 446–453. doi: 10/f4pwb3

Kushnir, T., Wellman, H. M., & Gelman, S. A. (2009). A self-agency bias in preschoolers' causal inferences. *Developmental Psychology*, *45*(2), 597. doi: 10/dnv723

Lancy, D. F. (2002). Cultural constraints on children's play. *Play and Culture Studies*, *4*, 53–62.

Lancy, D. F. (2007). Accounting for variability in mother–child play. *American Anthropologist*, *109*(2), 273–284. doi: 10/fs796m

Landrum, A., Mills, C., & Johnston, A. (2013). When do children trust the expert? Benevolence information influences children's trust more than expertise. *Developmental Science*, *16*, 622–638. doi: 10/f48466

Lapidow, E., Killeen, I., & Walker, C. M. (2022). Learning to recognize uncertainty vs. recognizing uncertainty to learn: Confidence judgments and exploration decisions in preschoolers. *Developmental Science*, *25*(2). doi: 10/gnj5xn

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

Laversanne-Finot, A., Péré, A., & Oudeyer, P.-Y. (2021). Intrinsically Motivated

Exploration of Learned Goal Spaces. *Frontiers in Neurorobotics*, *14*. doi: 10.3389/fnbot.2020.555271

Lawrence, A. (1987). Consumer demand theory and the assessment of animal welfare. *Animal Behavior*, *35*(1), 293–295. doi: 10.1016/s0003-3472(87)80236-1

Lee, D., Seo, H., & Jung, M. W. (2012). Neural Basis of Reinforcement Learning and Decision Making. *Annual Review of Neuroscience*, *35*(1), 287–308. doi: 10.1146/annurev-neuro-062111-150512

Lee, P. C. (1982). Play as a means for developing relationships. *Primate Social Relationships*.

Legare, C. (2014). The contributions of explanation and exploration to children's scientific reasoning. *Child Development Perspectives*, *8*, 101–106.

Legare, C., Gelman, S. A., & Wellman, H. (2010). Inconsistency with prior knowledge triggers children's causal explanatory reasoning. Child. *Child Development*, *81*, 929–944. doi: 10.1111/j.1467-8624.2010.01443.x

Legare, C., Mills, C., Souza, A., Plummer, L., & Yasskin, R. (2013). The use of questions as problem-solving strategies during early childhood. *Journal of Experimental Child Psychology*, *114*, 63–76. doi: 10/f4jghf

Legare, C., & Nielsen, M. (2015). Imitation and innovation: The dual engines of cultural learning. *Trends in Cognitive Sciences*, *19*(11), 688–699. doi: 10/f7zb9w

Legare, C., Wen, N. J., Herrmann, P. A., & Whitehouse, H. (2015). Imitative flexibility and the development of cultural learning. *Cognition*, *142*, 351–361. doi: 10/f7md9m

Legare, C. H. (2012). Exploring explanation: Explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child Development*, *83*(1), 173–185. doi: 10.1111/j.1467-8624.2011.01691.x

Leonard, J. A., Bennett-Pierre, G., & Gweon, H. (2019). Who is better? preschoolers infer relative competence based on efficiency of process and quality of outcome. In *Proceedings of the 41st Annual Conference of the Cognitive Science Society.*

Leonard, J. A., Garcia, A., & Schulz, L. E. (2020). How Adults' Actions, Outcomes, and Testimony Affect Preschoolers' Persistence. *Child Development*, *91*(4), 1254–1271. doi: 10/gjkznr

Leonard, J. A., Lee, Y., & Schulz, L. E. (2017). Infants make more attempts to achieve a goal when they see adults persist. *Science*, *357*(6357), 1290–1294.

Lewis, D. (1978). Truth in Fiction. *American Philosophical Quarterly*, *15*(1), 37–46.

Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, e1. doi: 10.1017/s0140525x1900061x

Lillard, A. S. (2001). Pretend play as twin earth: A social-cognitive analysis. *Developmental Review*, *21*(4), 495–531. doi: 10/cd8ngg

Lillard, A. S. (2012). Preschool children's development in classic montessori, supplemented montessori, and conventional programs. *Journal of School Psychology*, *50*(3), 379–401. doi: 10.1016/j.jsp.2012.01.001

Lillard, A. S. (2015). The Development of Play. In L. Liben & U. Mueller (Eds.), *Handbook of Child Psychology and Developmental Science, Vol. 3: Cognitive Development* (pp. 425–468). New York: Wiley-Blackwell. doi: https://doi.org/

10.1002/9781118963418.childpsy211
Lillard, A. S. (2017). Why Do the Children (Pretend) Play? *Trends in Cognitive Sciences*, *21*(11), 826–834. doi: 10/gcf9dh
Lillard, A. S., Lerner, M. D., Hopkins, E. J., Dore, R. A., Smith, E. D., & Palmquist, C. M. (2013). The impact of pretend play on children's development: A review of the evidence. *Psychological Bulletin*, *139*(1), 1–34. doi: 10.1037/a0029321
Lillard, A. S., Pinkham, A., & Smith, E. (2011). Pretend play and cognitive development. In *The Wiley-Blackwell handbook of childhood cognitive development* (2nd ed ed., pp. 285–311). Malden, MA: Wiley-Blackwell.
Lin, Y., & Westgate, E. C. (n.d.). The origins of boredom. In L. Al-Shawif & T. Shackelford (Eds.), *The Oxford Handbook of Evolution and the Emotions.* Oxford University Press.
Linke, C., Ady, N. M., White, M., Degris, T., & White, A. (2020). Adapting Behavior via Intrinsic Reward: A Survey and Empirical Study. *Journal of Artificial Intelligence Research*, *69*, 1287–1332. doi: 10/gq6jcg
Liquin, E. G., & Lombrozo, T. (2020a). Explanation-seeking curiosity in childhood. *Current Opinion in Behavioral Sciences*, *35*, 14–20. doi: 10/gmmkks
Liquin, E. G., & Lombrozo, T. (2020b). A functional approach to explanation-seeking curiosity. *Cognitive Psychology*, *119*, 101276. doi: 10/grwms6
Liquin, E. G., Metz, S., & Lombrozo, T. (2020). Science demands explanation, religion tolerates mystery. *Cognition*, *204*, 104398. doi: 10/gkr7h6
Little, D. Y., & Sommer, F. T. (2013). Learning and exploration in action-perception loops. *Frontiers in Neural Circuits*, *7*. doi: 10.3389/fncir.2013.00037
Liu, S., Gonzalez, G., & Warneken, F. (2019). Worth the wait: Children trade off delay and reward in self- and other-benefiting decisions. *Developmental Science*, *22*(1), e12702. doi: 10/ggbvwh
Liu, S., & Spelke, E. S. (2017). Six-month-old infants expect agents to minimize the cost of their actions. *Cognition*, *160*, 35–42. doi: 10.1016/j.cognition.2016.12.007
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. doi: 10.1126/science.aag2132
Lockman, J. J. (2000). A perception–action perspective on tool use development. *Child Development*, *71*(1), 137–144. doi: 10/fbd3gx
Lockman, J. J., & Tamis-LeMonda, C. S. (2021). Young Children's Interactions with Objects: Play as Practice and Practice as Play. *Annual Review of Developmental Psychology*, *3*(1), 165–186. doi: 10.1146/annurev-devpsych-050720-102538
Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological bulletin*, *116*(1), 75. doi: 10.1037/0033-2909.116.1.75
Lombrozo, T. (2011). Philosophy Compass. *The instrumental value of explanations*, *6*, 539–551.
Lombrozo, T. (2016). Explanatory Preferences Shape Learning and Inference. *Trends in Cognitive Sciences*, *20*(10), 748–759. doi: 10.1016/j.tics.2016.08.001
Lombrozo, T. (2019). 'Learning by thinking' in science and in everyday life. *The scientific imagination*, 230–249.
Lopes, M., Lang, T., Toussaint, M., & Oudeyer, P.-Y. (2012). Exploration in

model-based reinforcement learning by empirically estimating learning progress. *Advances in neural information processing systems*, *25*.
Loveland, K. A. (1987). Behavior of young children with down syndrome before the mirror: Finding things reflected. *Child Development*, 928–936. doi: 10/bvq6zs
Lubinski, D. (2000). Scientific and Social Significance of Assessing Individual Differences: "Sinking Shafts at a Few Critical Points". *Annual Review of Psychology*, *51*(1), 405–444. doi: 10/b7j67t
Lucca, K., Horton, R., & Sommerville, J. A. (2020). Infants rationally decide when and how to deploy effort. *Nature Human Behaviour*, *4*(4), 372–379. doi: 10/gq6jbq
Luchkina, E., Sommerville, J. A., & Sobel, D. M. (2018). More than just making it go: Toddlers effectively integrate causal efficacy and intentionality in selecting an appropriate causal intervention. *Cognitive Development*, *45*, 48–56. doi: 10.1016/j.cogdev.2017.12.003
Luo, Y. (2011). Three-month-old infants attribute goals to a non-human agent. *Developmental Science*, *14*(2), 453–460. doi: 10.1111/j.1467-7687.2010.00995.x
Lynch, C., Khansari, M., Xiao, T., Kumar, V., Tompson, J., Levine, S., & Sermanet, P. (2020). Learning latent plans from play. In L. P. Kaelbling, D. Kragic, & K. Sugiura (Eds.), *Proceedings of the conference on robot learning* (Vol. 100, pp. 1113–1132). Pmlr.
Lyons, D. E., Young, A. G., & Keil, F. C. (2007). The hidden structure of overimitation. *Proceedings of the National Academy of Sciences*, *104*(50), 19751–19756. doi: 10/fj7x9d
Ma, L., & Ganea, P. (2010). Dealing with conflicting information: Young children's reliance on what they see versus what they are told. *Developmental Science*, *13*, 151–160. doi: 10/dj52g8
Ma, W. J., & Woodford, M. (2020). Multiple conceptions of resource rationality. *Behavioral and Brain Sciences*, *43*.
MacDonald, K. (2007). Cross-cultural comparison of learning in human hunting. *Human Nature*, *18*(4), 386–402. doi: 10.1007/s12110-007-9019-8
MacDonald, K., Schug, M., Chase, E., & Barth, H. (2013). My people, right or wrong? Minimal group membership disrupts preschoolers' selective trust. *Cognitive Development*, *28*, 247–259. doi: 10/gfzkr9
Magid, R., DePascale, M., & Schulz, L. E. (2018). Four- and 5-year-olds infer differences in relative ability and appropriately allocate roles to achieve cooperative, competitive, and prosocial goals. *Open Mind*, *2*(2), 72–85. doi: 10.1162/opmi_a_00019
Magid, R., Sheskin, M., & Schulz, L. E. (2015). Imagination and the generation of new ideas. *Cognitive Development*, *34*, 99–110.
Marazita, J. M., & Merriman, W. E. (2004). Young children's judgment of whether they know names for objects: The metalinguistic ability it reflects and the processes it involves. *Journal of Memory and Language*, *51*(3), 458–472. doi: 10/d282zn
Marcovitch, S., Boseovski, J. J., & Knapp, R. J. (2007). Use it or lose it: Examining preschoolers' difficulty in maintaining and executing a goal. *Developmental Science*, *10*(5), 559–564.

Martin, P., & Caro, T. M. (1985). On the functions of play and its role in behavioral development. *Advances in the Study of Behavior*, *15*, 59–103. doi: 10.1016/s0065-3454(08)60487-8

Martius, G., Der, R., & Ay, N. (2013). Information driven self-organization of complex robotic behaviors. *PLoS ONE*, *8*(5), e63400. doi: 10.1371/journal.pone.0063400

Mascaro, O., & Csibra, G. (2022). Infants expect agents to minimize the collective cost of collaborative actions. *Scientific Reports*, *12*(1), 17088. doi: 10/gq6jcq

Mather, J. A., & Anderson, R. C. (1999). Exploration, play and habituation in octopuses (octopus dofleini). *Journal of Comparative Psychology*, *113*(3), 333–338. doi: 10.1037/0735-7036.113.3.333

Maye, J., Werker, J. F., & Gerken, L. (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition*, *82*(3), B101–b111. doi: 10.1016/s0010-0277(01)00157-3

McCall, R. B., & Carriger, M. S. (1993). A meta-analysis of infant habituation and recognition memory performance as predictors of later iq. *Child Development*, *64*(1), 57–79. doi: 10/d9hqct

McCarthy, A. M., Schoorman, F. D., & Cooper, A. C. (1993). Reinvestment decisions by entrepreneurs: rational decision-making or escalation of commitment? *Journal of Business Venturing*, *8*(1), 9–24.

McCarty, M. E., Clifton, R. K., & Collard, R. R. (1999). Problem solving in infancy: The emergence of an action plan. *Developmental Psychology*, *35*(4), 1091. doi: 10.1037/0012-1649.35.4.1091

McCarty, M. E., Clifton, R. K., & Collard, R. R. (2001). The beginnings of tool use by infants and toddlers. *Infancy*, *2*(2), 233–256. doi: 10.1207/s15327078in0202_8

McCormack, T., Frosch, C., Patrick, F., & Lagnado, D. (2015). Temporal and statistical information in causal structure learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(2), 395–416. doi: 10.1037/a0038385

McCowan, B., Marino, L., Vance, E., Walke, L., & Reiss, D. (2000). Bubble ring play of bottlenose dolphins (tursiops truncatus): Implications for cognition. *Journal of Comparative Psychology*, *114*(1), 98. doi: 10/fq7v3w

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

Mears, C. E., & Harlow, H. F. (1975). Play: early and eternal. *Proceedings of the National Academy of Sciences*, *72*(5), 1878–1882. doi: 10.1073/pnas.72.5.1878

Meder, B., Wu, C. M., Schulz, E., & Ruggeri, A. (2021). Development of directed and random exploration in children. *Developmental Science*, *24*(4). doi: 10/gq6jcr

Mercier, H., Bernard, S., & Clement, F. (2014). Early sensitivity to arguments: how preschoolers weight circular arguments. *Journal of Experimental Child Psychology*, *125*, 102–109. doi: 10.1016/j.jecp.2013.11.011

Mills, C. (2013). Knowing when to doubt: Developing a critical stance when learning from others. *Developmental Psychology*, *49*, 404–418. doi: 10/f4p544

Mills, C. M., Danovitch, J. H., Rowles, S. P., & Campbell, I. L. (2017). Children's success at detecting circular explanations and their interest in future learning. *Psychonomic Bulletin & Review*, *24*(5), 1465–1477. doi: 10/gk8qb2

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. doi: 10.1080/15248372.2010.516419

Mills, C. M., Legare, C. H., Grant, M. G., & Landrum, A. R. (2011). Determining who to question, what to ask, and how much information to ask for: The development of inquiry in young children. *Journal of Experimental Child Psychology*, *110*(4), 539-560. doi: 10.1016/j.jecp.2011.06.003

Montessori, M. (1912/1964). *The Montessori Method*. Ny: Shocken.

Morgan, T., Laland, K., & Harris, P. (2014). The development of adaptive conformity in young children: Effects of uncertainty and consensus. *Developmental Science*, *18*, 511–524. doi: 10/f7d68k

Muentener, P., Herrig, E., & Schulz, L. E. (2018). The efficiency of infants' exploratory play is related to longer-term cognitive development. *Frontiers in Psychology*, *9*.

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

Mulcahy, N. J., & Call, J. (2006). Apes save tools for future use. *science*, *312*(5776), 1038–1040. doi: 10/bzxtrc

Munakata, Y., Snyder, H. R., & Chatham, C. H. (2012). Developing cognitive control: Three key transitions. *Current directions in psychological science*, *21*(2), 71–77.

Murray, L., & Trevarthen, C. (1986). The infant's role in mother–infant communications. *Journal of Child Language*, *13*(1), 15–29. doi: 10.1017/s0305000900000271

Naiman, R. J., Johnston, C. A., & Kelley, J. C. (1988). Alteration of North American streams by beaver. *BioScience*, *38*(11), 753–762. doi: 10.2307/1310784

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). doi: 10.1016/j.cognition.2013.09.007

Nguyen, S. (2012). The role of external sources of information in children's evaluative food categories. *Infant and Child Development*, *21*, 216–235. doi: 10/bvfjhn

Nielsen, M., Cucchiaro, J., & Mohamedally, J. (2012). When the Transmission of Culture Is Child's Play. *Plos One*, *7*(3), e34066. doi: 10/gq6jcm

Nielsen, M., Moore, C., & Mohamedally, J. (2012). Young children overimitate in third-party contexts. *Journal of Experimental Child Psychology*, *112*(1), 73–83. doi: 10/fzzs8t

Nishida, T., & Hiraiwa, M. (1982). Natural history of a tool-using behavior by wild chimpanzees in feeding upon wood-boring ants. *Journal of Human Evolution*, *11*(1), 73–99. doi: 10.1016/s0047-2484(82)80033-x

Nissen, H. W. (1930). A study of exploratory behavior in the white rat by means of the obstruction method. *The Pedagogical Seminary and Journal of Genetic Psychology*, *37*(3), 361–376.

Nurmsoo, E., & Robinson, E. (2009). Children's trust in previously inaccurate informants who were well or poorly informed: When past errors can be excused. *Child Development*, *80*, 23–27. doi: 10/crsj42

Nwokah, E. E., & Ikekeonwu, C. (1998). A sociocultural comparison of nigerian and

american children's games. *Play and Culture*, *1*, 59–76.
O'Hare, D., & Smitheram, T. (1995). 'Pressing on'into deteriorating conditions: An application of behavioral decision theory to pilot decision making. *The International Journal of Aviation Psychology*, *5*(4), 351–370. doi: 10.1207/s15327108ijap0504_2
Okonofua, E. C., Simpson, K. N., Jesri, A., Rehman, S. U., Durkalski, V. L., & Egan, B. M. (2006). Therapeutic inertia is an impediment to achieving the healthy people 2010 blood pressure control goals. *Hypertension*, *47*(3), 345–351.
Onishi, K. H., & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, *308*(5719), 255–258. doi: 10/bj7ndt
Ossmy, O., Hoch, J. E., MacAlpine, P., Hasan, S., Stone, P., & Adolph, K. E. (2018). Variety Wins: Soccer-Playing Robots and Infant Walking. *Frontiers in Neurorobotics*, *12*. doi: 10.3389/fnbot.2018.00019
Oudeyer, P.-Y. (2018). Computational Theories of Curiosity-Driven Learning. *arXiv:1802.10546 [cs]*.
Oudeyer, P.-Y., Gottlieb, J., & Lopes, M. (2016). Intrinsic motivation, curiosity, and learning: Theory and applications in educational technologies. In *Progress in brain research* (Vol. 229, pp. 257–284). Elsevier.
Oudeyer, P.-Y., & Kaplan, F. (2009). What is intrinsic motivation? A typology of computational approaches. *Frontiers in Neurorobotics*, *1*, 1–14. doi: 10.3389/neuro.12.006.2007
Oudeyer, P.-Y., Kaplan, F., & Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE transactions on evolutionary computation*, *11*(2), 265–286. doi: 10.1109/tevc.2006.890271
Over, H., & Carpenter, M. (2013). The Social Side of Imitation. *Child Development Perspectives*, *7*(1), 6–11. doi: 10.1111/cdep.12006
Pace, D. S. (2000). Fluke-made bubble rings as toys in bottlenose dolphin calves (tursiops truncatus). *Aquatic Mammals*, *26*(1), 57–64.
Palagi, E. (2006). Social play in bonobos (pan paniscus) and chimpanzees (pan troglodytes): Implicationsfor natural social systems and interindividual relationships. *American Journal of Physical Anthropology*, *129*(3), 418–426.
Palagi, E. (2008). Sharing the motivation to play: the use of signals in adult bonobos. *Animal Behaviour*, *75*(3), 887–896. doi: 10.1016/j.anbehav.2007.07.016
Panksepp, J. (1981). The ontogeny of play in rats. *Developmental Psychobiology: The Journal of the International Society for Developmental Psychobiology*, *14*(4), 327–332. doi: 10.1002/dev.420140405
Panksepp, J., & Burgdorf, J. (2010). Laughing rats? playful tickling arouses high-frequency ultrasonic chirping in young rodents. *American Journal of Play*, *2*(3), 357–372.
Pasquini, E. S., Corriveau, K. H., Koenig, M. A., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology*, *43*(5), 1216–1226. doi: 10.1037/0012-1649.43.5.1216
Pathak, D., Agrawal, P., Efros, A. A., & Darrell, T. (2017). Curiosity-Driven Exploration by Self-Supervised Prediction. In *Proceedings of the 34th International Conference on Machine Learning* (pp. 2778–2787).

Peirce, C. S. (1878). How to make our ideas clear. *Popular Science Monthly*, *12*, 286–302.

Pellegrini, A. D., Dupuis, D., & Smith, P. K. (2007). Play in evolution and development. *Developmental Review*, *27*(2), 261–276. doi: 10.1016/j.dr.2006.09.001

Pellegrini, A. D., & Smith, P. K. (1998). Physical Activity Play: The Nature and Function of a Neglected Aspect of Play. *Child Development*, *69*(3), 577–598. doi: 10/bkr2hh

Pellis, S. M., & Iwaniuk, A. N. (1999). The roles of phylogeny and sociality in the evolution of social play in muroid rodents. *Animal Behaviour*, *58*(2), 361–373. doi: 10.1006/anbe.1999.1141

Pellis, S. M., & Iwaniuk, A. N. (2000). Adult–adult play in primates: comparative analyses of its origin, distribution and evolution. *Ethology*, *106*(12), 1083–1104. doi: 10/fn7rd9

Pellis, S. M., & Pellis, V. C. (1991). Role reversal changes during the ontogeny of play fighting in male rats: Attack vs. defense. *Aggressive Behavior*, *17*(3), 179–189. doi: 10.1002/1098-2337(1991)17:3<179::aid-ab2480170306>3.0.co;2-q

Pelz, M., & Kidd, C. (2020). The elaboration of exploratory play. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *375*(1803), 20190503. doi: 10.1098/rstb.2019.0503

Pelz, M., Schulz, L., & Jara-Ettinger, J. (2020). *The Signature of All Things: Children Infer Knowledge States from Static Images* (preprint). PsyArXiv. doi: 10.31234/osf.io/f692k

Perez, J., & Feigenson, L. (2022). Violations of expectation trigger infants to search for explanations. *Cognition*, *218*, 104942. doi: 10.31234/osf.io/eahjd

Pesowski, M. L., Quy, A. D., Lee, M., & Schachner, A. (2020). Children use inverse planning to detect social transmission in design of artifacts. *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*.

Petrů, M., Špinka, M., Charvátová, V., & Lhota, S. (2009). Revisiting play elements and self-handicapping in play: A comparative ethogram of five Old World monkey species. *Journal of Comparative Psychology*, *123*(3), 250–263. doi: 10.1037/a0016217

Phillips, J., Morris, A., & Cushman, F. (2019). How We Know What Not To Think. *Trends in Cognitive Sciences*, *23*(12), 1026–1040. doi: 10.1016/j.tics.2019.09.007

Phillips, L. S., Branch Jr, W. T., Cook, C. B., Doyle, J. P., El-Kebbi, I. M., Gallina, D. L., ... Barnes, C. S. (2001). Clinical inertia. *Annals of Internal Medicine*, *135*(9), 825–834.

Piaget, J. (1926). *The Language and Thought of the Child.* Harcourt, Brace.

Piaget, J. (1929). *The Child's Conception of the World.* Routledge.

Piaget, J. (1952). *The origins of intelligence in children. (M. Cook, Trans.).* New York, NY: WW Norton.

Piaget, J. (1954). *The construction of reality in the child.* Ny: Basic Books.

Piaget, J. (1962). Play. In *Dreams, and imitation in childhood.* New York, NY: WW Norton.

Plötner, M., Over, H., Carpenter, M., & Tomasello, M. (2015). The effects of collaboration and minimal-group membership on children's prosocial behavior, liking, affiliation, and trust. *Journal of Experimental Child Psychology*, *139*, 161–173. doi: 10/gh3skk

Poirier, F. E. (1969). The nilgiri langur (presbytis johnii) troop: its composition, structure, function and change. *Folia Primatologica*, *10*(1-2), 20–47.

Poirier, F. E., & Smith, E. O. (1974). Socializing functions of primate play. *American Zoologist*, *14*(1), 275–287. doi: 10.1093/icb/14.1.275

Pope, D. G., & Schweitzer, M. E. (2011). Is tiger woods loss averse? persistent bias in the face of experience, competition, and high stakes. *American Economic Review*, *101*(1), 129–57.

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. doi: 10.1073/pnas.1304326110

Premack, D. (2004). Is Language the Key to Human Intelligence? *Science*, *303*(5656), 318–320. doi: 10.1126/science.1093993

Prentice, N., & Gordon, D. (1987). Santa Claus and the Tooth Fairy for the Jewish child and parent. *The Journal of Genetic Psychology*, *148*, 139–151. doi: 10/cg9frx

Prentice, N., Manosevitz, M., & Hubbs, L. (1978). Imaginary figures of early childhood: Santa claus, easter bunny, and the tooth fairy. *American Journal of Orthopsychiatry*, *48*(4), 618. doi: 10/ck6734

Quinn, S., Donnelly, S., & Kidd, E. (2018). The relationship between symbolic play and language acquisition: A meta-analytic review. *Developmental Review*, *49*, 121–135. doi: 10/gd9jdj

Raby, C. R., Alexis, D. M., Dickinson, A., & Clayton, N. S. (2007). Planning for the future by western scrub-jays. *Nature*, *445*(7130), 919–921. doi: 10/dk2q2v

Raine, A., Reynolds, C., Venables, P. H., & Mednick, S. A. (2002). Stimulation seeking and intelligence: a prospective longitudinal study. *Journal of Personality and Social Psychology*, *82*(4), 663. doi: 10/ccxjtp

Rakoczy, H., Warneken, F., & Tomasello, M. (2008). The sources of normativity: Young children's awareness of the normative structure of games. *Developmental Psychology*, *44*(3), 875–881. doi: 10/br6p85

Raz, G., & Saxe, R. (2020). Learning in Infancy Is Active, Endogenously Motivated, and Depends on the Prefrontal Cortices. *Annual Review of Developmental Psychology*, *2*(1), 247–268. doi: 10.1146/annurev-devpsych-121318-084841

Réale, D., Reader, S. M., Sol, D., McDougall, P. T., & Dingemanse, N. J. (2007). Integrating animal temperament within ecology and evolution. *Biological reviews*, *82*(2), 291–318. doi: 10/bh3tz4

Redshaw, J., Taylor, A. H., & Suddendorf, T. (2017). Flexible planning in ravens? *Trends in Cognitive Sciences*, *21*(11), 821–822. doi: 10/grwmn6

Reppert, S. M., Gegear, R. J., & Merlin, C. (2010). Navigational mechanisms of migrating monarch butterflies. *Trends in neurosciences*, *33*(9), 399–406. doi: 10.1016/j.tins.2010.04.004

Reyes-Jaquez, B., & Echols, C. (2013). Developmental differences in the relative weighing of informants' social attributes. *Developmental Psychology*, *49*, 602–613. doi: 10/f4p6bq

Richardson, E., & Keil, F. C. (2022). Thinking takes time: Children use agents' response times to infer the source, quality, and complexity of their knowledge. *Cognition*, *224*, 105073.

Roberts, J. M., & Sutton-Smith, B. (1962). Child training and game involvement. *Ethnology*, *1*(2), 166–185. doi: 10/dv6ct8

Rochat, P. (1989). Object Manipulation and Exploration in 2- to 5-Month-Old Infants. *Developmental Psychology*, *25*(6), 871–884. doi: 10.1037/0012-1649.25.6.871

Ronfard, S., Zambrana, I., Hermansen, T., & Kelemen, D. (2018). Question-asking in childhood: A review of the literature and a framework for understanding its development. *Developmental Review*, *49*, 101–120. doi: 10/gd9fvz

Royka, A., Chen, A., Aboody, R., Huanca, T., & Jara-Ettinger, J. (2022). People infer communicative action through an expectation for efficient communication. *Nature Communications*, *13*(1), 4160. doi: 10/gq6jcn

Ruba, A. L., & Repacholi, B. M. (2020). Do Preverbal Infants Understand Discrete Facial Expressions of Emotion? *Emotion Review*, *12*(4), 235–250. doi: 10.1177/1754073919871098

Ruff, H. A., McCarton, C., Kurtzberg, D., & Vaughan Jr, H. G. (1984). Preterm infants' manipulative exploration of objects. *Child Development*, 1166–1173. doi: 10/cxn597

Ruggeri, A., & Lombrozo, T. (2015). Children adapt their questions to achieve efficient search. *Cognition*, *143*(3), 203–216. doi: 10/gq6jbp

Ruggeri, A., Lombrozo, T., Griffiths, T. L., & Xu, F. (2016). Sources of developmental change in the efficiency of information search. *Developmental Psychology*, *52*(12), 2159–2173. doi: 10/f9gbx8

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), 1620–1632. doi: 10/gcjpjg

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. doi: 10.1016/j.cognition.2019.104013

Rutz, C., Bluff, L. A., Reed, N., Troscianko, J., Newton, J., Inger, R., ... Bearhop, S. (2010). The ecological significance of tool use in new caledonian crows. *Science*, *329*(5998), 1523–1526. doi: 10.1126/science.1192053

Ruvolo, P., Messinger, D., & Movellan, J. (2015). Infants Time Their Smiles to Make Their Moms Smile. *Plos One*, *10*(9), e0136492. doi: 10.1371/journal.pone.0136492

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. doi: 10/bgj5kz

Sabbagh, M., & Baldwin, D. (2001). Learning words from knowledgeable versus ignorant speakers: Links between preschoolers' theory of mind and semantic development. *Child Development*, *72*, 1054–1070. doi: 10/d6wt8k

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical Learning by 8-Month-Old Infants. *Science*, *274*(5294), 1926–1928. doi: 10/fcqz9d

Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, *70*(1), 27–52. doi: 10/bjxdg9

Salge, C., Glackin, C., & Polani, D. (2013). Empowerment - an introduction. *CoRR*, *abs/1310.1863*.

Sato, N., Tan, L., Tate, K., & Okada, M. (2015). Rats demonstrate helping behavior toward a soaked conspecific. *Animal Cognition*, *18*(5), 1039–1047. doi: 10.1007/s10071-015-0872-2

Savickas, M. L., & Spokane, A. R. (Eds.). (1999). *Vocational interests: Meaning, measurement, and counseling use*. Palo Alto, CA, US: Davies-Black Publishing.

Saylor, M. M., Baldwin, D. A., Baird, J. A., & LaBounty, J. (2007). Infants' on-line segmentation of dynamic human action. *Journal of Cognition and Development*, *8*(1), 113–128. doi: 10.1080/15248370709336996

Sazima, I. (2008). Playful birds: cormorants and herons play with objects and practice their skills. *Biota Neotropica*, *8*(2). doi: 10/dkjqww

Scarlett, W. G., Naudeau, S., Salonius-Pasternak, D., & Ponte, I. (2005). *Children's Play*. Sage.

Schachner, A., & Carey, S. (2013). Reasoning about 'irrational' actions: When intentional movements cannot be explained, the movements themselves are seen as the goal. *Cognition*, *129*(2), 309–327.

Schacter, D. L. (2001). *The seven sins of memory: How the mind forgets and remembers*. Boston, MA, US: Houghton, Mifflin and Company.

Schleihauf, H., Graetz, S., Pauen, S., & Hoehl, S. (2018). Contrasting social and cognitive accounts on overimitation: The role of causal transparency and prior experiences. *Child Development*, *89*(3), 1039–1055. doi: 10/gdkgsg

Schmidhuber, J. (1991). Curious model-building control systems. *IEEE IJCNN*, 1458–1463.

Schmidhuber, J. (2010). Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, *2*(3), 230–247. doi: 10.1109/tamd.2010.2056368

Schmidhuber, J. (2013). Powerplay: Training an increasingly general problem solver by continually searching for the simplest still unsolvable problem. *Frontiers in Psychology*, *4*. doi: 10.3389/fpsyg.2013.00313

Schulz, L. E. (2012a). Finding new facts; thinking new thoughts. In *Rational constructivism in cognitive development* (pp. 269–294). Elsevier.

Schulz, L. E. (2012b). The origins of inquiry: Inductive inference and exploration in early childhood. *Trends in Cognitive Sciences*, *16*(7), 382–389.

Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: Preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, *43*(4), 1045–1050.

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

Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science*, *10*(3),

322–332. doi: 10.1111/j.1467-7687.2007.00587.x
Schulz, L. E., Hooppell, K., & Jenkins, A. (2008). Judicious imitation: Young children imitate deterministic actions exactly, stochastic actions more variably. *Child Development*, *79*(20), 395–410. doi: 10.1111/j.1467-8624.2007.01132.x
Schulz, L. E., & Sommerville, J. (2006). God does not play dice: Causal determinism and preschoolers? causal inferences. *Child Development*, *77*, 427–42.
Schulz, L. E., Standing, H. R., & Bonawitz, E. B. (2008). Word, thought, and deed: the role of object categories in children's inductive inferences and exploratory play. *Developmental Psychology*, *44*(5), 1266.
Schwartzman, H. B. (1986). A cross-cultural perspective on child-structured play activities and materials. *Play interactions: The contribution of play materials and parental involvement to children's development*, 13–30.
Seeley, T. D. (2014). *Honeybee ecology: a study of adaptation in social life* (Vol. 36). Princeton University Press.
Senju, A., Southgate, V., Snape, C., Leonard, M., & Csibra, G. (2011). Do 18-month-olds really attribute mental states to others? A critical test. *Psychological Science*, *22*, 878–880. doi: 10/dbx7q8
Shafto, P., Eaves, B., Navarro, D., & Perfors, A. (2012). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental Science*, *15*, 436–447. doi: 10/grf3b6
Sharon, T., & Woolley, J. (2004). Do monsters dream? Young children's understanding of the fantasy/reality distinction. *British Journal of Developmental Psychology*, *22*, 293–310. doi: 10/brc4j8
Sharpe, F. A., & Dill, L. M. (1997). The behavior of pacific herring schools in response to artificial humpback whale bubbles. *Canadian Journal of Zoology*, *75*, 725–730.
Sharpe, L. L. (2005). Play does not enhance social cohesion in a cooperative mammal. *Animal Behaviour*, *70*(3), 551–558. doi: 10.1016/j.anbehav.2004.08.025
Sharpe, L. L., & Cherry, M. I. (2003). Social play does not reduce aggression in wild meerkats. *Animal Behaviour*, *66*(5), 989–997. doi: 10/fq395d
Sharpe, L. L., Clutton-Brock, T. H., Brotherton, P. N., Cameron, E. Z., & Cherry, M. I. (2002). Experimental provisioning increases play in free-ranging meerkats. *Animal Behaviour*, *64*(1), 113–121.
Shneidman, L., Gweon, H., Schulz, L. E., & Woodward, A. L. (2016). Learning from others and spontaneous exploration: A cross-cultural investigation. *Child Development*, *87*(3), 723–735. doi: 10.1111/cdev.12502
Siegel, M. H., Magid, R. W., Pelz, M., Tenenbaum, J. B., & Schulz, L. E. (2021). Children's exploratory play tracks the discriminability of hypotheses. *Nature Communications*, *12*(1), 3598. doi: 10.1038/s41467-021-23431-2
Sigman, M. (1976). Early development of preterm and full-term infants: Exploratory behavior in eight-month-olds. *Child Development*, 606–612. doi: 10/c72497
Silvia, P. J. (2012). Curiosity and motivation. *The Oxford handbook of human motivation*, 157–166. doi: 10/gq6jb9
Sim, Z. L., Mahal, K., & Xu, F. (2017). Learning about causal systems through play. In *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. doi: 10/gq6jb6

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. doi: 10/f94gdw

Sim, Z. L., & Xu, F. (2019). Another Look at Looking Time: Surprise as Rational Statistical Inference. *Topics in Cognitive Science*, *11*(1), 154–163. doi: 10/gn8z6p

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*(1), 99–118.

Simon, H. A., & Newell, A. (1971). Human problem solving: The state of the theory in 1970. *American Psychologist*.

Singer, D. G., Singer, J. L., D'Agnostino, H., & DeLong, R. (2009). Children's pastimes and play in sixteen nations: Is free-play declining?. *American Journal of Play*, *1*(3), 283–312.

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

Slater, A., & Morison, V. (1991). Visual attention and memory at birth. In *Newborn attention: Biological constraints and the influence of experience* (pp. 256–277). Westport, CT, US: Ablex Publishing.

Smith, E. F. (1991). The influence of nutrition and postpartum mating on weaning and subsequent play behaviour of hooded rats. *Animal Behaviour*, *41*(3), 513–524. doi: 10.1016/s0003-3472(05)80854-1

Smith, P. K. (1982). Does play matter? functional and evolutionary aspects of animal and human play. *Behavioral and Brain Sciences*, *5*(1), 139–155. doi: 10.1017/s0140525x0001092x

Sobel, D., & Corriveau, K. (2010). Children monitor individuals' expertise for word learning. *Child Development*, *81*, 669–679. doi: 10/c2xgt7

Sobel, D., & Kushnir, T. (2013). Knowledge matters: How children evaluate the reliability of testimony as a process of rational inference. *Psychological Review*, *120*, 779–797. doi: 10/f5grwb

Sobel, D., & Macris, D. (2013). Children's understanding of speaker reliability between lexical and syntactic knowledge. *Developmental Psychology*, *49*, 523–532. doi: 10/f4pwp3

Sobel, D. M., Benton, D., Finiasz, Z., Taylor, Y., & Weisberg, D. S. (2022). The influence of children's first action when learning causal structure from exploratory play. *Cognitive Development*, *63*, 101194. doi: 10/gq6jcs

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

Sobel, D. M., Yoachim, C. M., Gopnik, A., Meltzoff, A. N., & Blumenthal, E. J. (2007). The blicket within: Preschoolers' inferences about insides and causes. *Journal of Cognition and Development*, *8*(2), 159–182. doi: 10.1080/15248370701202356

Sodian, B., Zaitchik, D., & Carey, S. (1991). Young children's differentiation of hypothetical beliefs from evidence. *Child Dev*, *62*(4), 753. doi: 10.2307/1131175

Sommerville, J. A., Enright, E. A., Horton, R. O., Lucca, K., Sitch, M. J., & Kirchner-Adelhart, S. (2018). Infants' prosocial behavior is governed by cost-benefit analyses. *Cognition*, *177*, 12–20. doi: 10.1016/j.cognition.2018.03.021

Sosa, F. A., & Ullman, T. (2022). *Type theory in human-like learning and inference.*

Spelke, E. S. (1985). Preferential-looking methods as tools for the study of cognition in infancy. In (pp. 323–363). Ablex Publishing.

Spelke, E. S. (2022). *What babies know: Core knowledge and composition volume 1* (Vol. 1). Oxford University Press.

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

Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*, *25*, 359–393. doi: 10/dmkgvf

Spinka, M., Newberry, R. C., & Bekoff, M. (2001). Mammalian play: training for the unexpected. *The Quarterly review of biology*, *76*(2), 141–168. doi: 10/d5625k

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

Stahl, A. E., & Feigenson, L. (2018). Violations of core knowledge shape early learning. *Topics in Cognitive Science*, *11*(1), 136–153. doi: 10.1111/tops.12389

Staw, B. M. (1976). Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action. *Organizational Behavior and Human Performance*, *16*(1), 27–44.

Stephens, E., Suarez, S., & Koenig, M. A. (2015). Early Testimonial Learning. In *Advances in Child Development and Behavior* (pp. 151–183).

Sukhbaatar, S., Kostrikov, I., Szlam, A., & Fergus, R. (2017). Intrinsic motivation and automatic curricula via asymmetric self-play. *CoRR*, *abs/1703.05407*.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268*(5217), 1632–1634.

Taylor, M., Cartwright, B., & Carlson, S. (1993). A developmental investigation of children's imaginary companions. *Developmental Psychology*, *29*, 276. doi: 10/cnjzkq

Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, *10*(7), 309–318. doi: 10.1016/j.tics.2006.05.009

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, *331*(6022), 1279–1285. doi: 10.1126/science.1192788

Tenney, E., Small, J., Kondrad, R., Jaswal, V. K., & Spellman, B. (2011). Accuracy, confidence, and calibration: How young children and adults assess credibility. *Developmental Psychology*, *47*, 1065–1077. doi: 10/bxr9dt

Thomas, H. (1965). Visual-fixation responses of infants to stimuli of varying complexity. *Child Development*, 629–638. doi: 10.2307/1126909

Thompson, K. V. (1998). Self assessment in juvenile play. In M. Bekoff & J. A. Byers

(Eds.), *Animal play: Evolutionary, comparative, and ecological perspectives* (pp. 183–204). Cambridge University Press.

Thornton, P., & Waterman-Pearson, A. (2002). Behavioural responses to castration in lambs. *Animal Welfare*, *11*(2), 203–212.

Thornton, S. (1995). *Children Solving Problems* (J. Bruner, M. Cole, & A. Karmiloff-Smith, Eds.).

Tizard, B., & Hughes, M. (1984). *Children learning at home and in school*. Fontana.

Todd, P. M., & Gigerenzer, G. (2012). *Ecological rationality: Intelligence in the world*. Oxford University Press. doi: 10.1093/acprof:oso/9780195315448.001.0001

Tomasello, M. (2022). *The evolution of agency: behavioral organization from lizards to humans*. Cambridge, Massachusetts: The MIT Press.

Tomasello, M., & Carpenter, M. (2007). Shared intentionality. *Developmental science*, *10*(1), 121–125. doi: 10/c6n78h

Tomasello, M., & Rakoczy, H. (2003). What Makes Human Cognition Unique? From Individual to Shared to Collective Intentionality. *Mind and Language*, *18*(2), 121–147. doi: 10.1111/1468-0017.00217

Traut, H. J., Chevalier, N., Guild, R. M., & Munakata, Y. (2021). Understanding and supporting inhibitory control: Unique contributions from proactive monitoring and motoric stopping to children's improvements with practice. *Child Development*, *92*(6). doi: 10.1111/cdev.13614

Tsividis, P., Tenenbaum, J., & Schulz, L. E. (2015). Hypothesis-space constraints in causal learning. In *Proceedings of the 37th annual conference of the cognitive science society* (pp. 2434–2439). Pasadena CA.

Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *Science*, *185*(4157), 1124–1131.

Uller, C. (2004). Disposition to recognize goals in infant chimpanzees. *Animal Cognition*, *7*, 154–161. doi: 10/d5qz6d

Ullman, T. D., Goodman, N. D., & Tenenbaum, J. B. (2012). Theory learning as stochastic search in the language of thought. *Cognitive Development*, *27*(4), 455–480.

Ullman, T. D., Spelke, E., Battaglia, P., & Tenenbaum, J. B. (2017). Mind Games: Game Engines as an Architecture for Intuitive Physics. *Trends in Cognitive Sciences*, *21*(9). doi: 10.1016/j.tics.2017.05.012

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*, 61–71.

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

Vasconcelos, M., Monteiro, T., & Kacelnik, A. (2015). Irrational choice and the value of information. *Scientific Reports*, *5*(1), 1–12. doi: 10/f7qc2r

Vergassola, M., Villermaux, E., & Shraiman, B. I. (2007). 'Infotaxis' as a strategy for searching without gradients. *Nature*, *445*(7126), 406–409. doi: 10/bc6ntn

Viholainen, H., Ahonen, T., Lyytinen, P., Cantell, M., LicSSc, A. T., & Lyytinen, H. (2006). Early motor development and later language and reading skills in children at risk of familial dyslexia. *Developmental Medicine & Child Neurology*, *48*(5), 367–373. doi: 10.1017/s001216220600079x

Vygotsky, L. (1934/1962). *Thought and language*. Cambridge, MA: MIT Press.

Walker, C. M., Bonawitz, E. B., & Lombrozo, T. (2017). Effects of explaining on children's preference for simpler hypotheses. *Psychonomic Bulletin & Review*, *24*(5), 1538–1547. doi: 10.3758/s13423-016-1144-0

Walker, C. M., & Gopnik, A. (2014). Toddlers Infer Higher-Order Relational Principles in Causal Learning. *Psychological Science*, *25*(1), 161–169. doi: 10/f5pfgf

Walker, C. M., Gopnik, A., & Ganea, P. (2015). Learning to learn from stories: Children's developing sensitivity to the causal structure of fictional worlds. *Child Development*, *86*, 310–318. doi: 10/f633t3

Walker, C. M., Lombrozo, T., Legare, C. H., & Gopnik, A. (2014). Explaining prompts children to privilege inductively rich properties. *Cognition*, *133*(2), 343–357. doi: 10.1016/j.cognition.2014.07.008

Walker, C. M., Lombrozo, T., Williams, J. J., Rafferty, A. N., & Gopnik, A. (2016). Explaining Constrains Causal Learning in Childhood. *Child Development*, *00*(0), 1–18. doi: 10.1111/cdev.12590

Watanabe, H. (1975). *Bow and arrow census in a west papuan lowland community: A new field for functional-ecological study*. Anthropology Museum, University of Queensland.

Watson, D. M. (1993). The play associations of red-necked wallabies (macropus rufogriseus banksianus) and relation to other social contexts. *Ethology*, *94*(1), 1–20. doi: 10.1111/j.1439-0310.1993.tb00543.x

Waytz, A., Dungan, J., & Young, L. (2013). The whistleblower's dilemma and the fairness–loyalty tradeoff. *Journal of Experimental Social Psychology*, *49*(6), 1027–1033.

Weisberg, D. S. (2015). Pretend play. *WIREs Cognitive Science*, *6*(3), 249–261. doi: 10/gd7fk9

Weisberg, D. S., & Bloom, P. (2009). Young children separate multiple pretend worlds. *Developmental Science*, *12*(5), 699–705. doi: 10.1111/j.1467-7687.2009.00819.x

Weisberg, D. S., & Gopnik, A. (2013). Pretense, counterfactuals, and bayesian causal models: Why what is not real really matters. *Cognitive Science*, *37*(7), 1368–1381. doi: 10.1111/cogs.12069

Weisberg, D. S., Hirsh-Pasek, K., & Golinkoff, R. M. (2013). Guided play: Where curricular goals meet a playful pedagogy. *Mind, Brain, and Education*, *7*(2), 104–112. doi: 10.1111/mbe.12015

Welker, W. I. (1956). Some determinants of play and exploration in chimpanzees. *Journal of Comparative and Physiological Psychology*, *49*(1), 84–89. doi: 10.1037/h0044463

Whiteley, P. F. (1995). Rational choice and political participation–evaluating the debate. *Political Research Quarterly*, *48*(1), 211–233.

Willatts, P. (1999). Development of means–end behavior in young infants: Pulling a support to retrieve a distant object. *Developmental Psychology*, *35*(3), 651–667.

doi: 10.1037/0012-1649.35.3.651
Wilson, K. P., Carter, M. W., Wiener, H. L., DeRamus, M. L., Bulluck, J. C., Watson, L. R., … Baranek, G. T. (2017). Object play in infants with autism spectrum disorder: A longitudinal retrospective video analysis. *Autism & developmental language impairments*, *2*, 2396941517713186. doi: 10.1177/2396941517713186
Wilson, T. D., Reinhard, D. A., Westgate, E. C., Gilbert, D. T., Ellerbeck, N., Hahn, C., … Shaked, A. (2014). Just think: The challenges of the disengaged mind. *Science*, *345*(6192), 75–77. doi: 10.1126/science.1250830
Wimpenny, J. H., Weir, A. A. S., & Kacelnik, A. (2010). New caledonian crows use tools for non-foraging activities. *Animal Cognition*, *14*(3), 459–464. doi: 10.1007/s10071-010-0366-1
Wing, L. A. (1995). Play is not the work of the child: Young children's perceptions of work and play. *Early Childhood Research Quarterly*, *10*(2), 223–247. doi: 10.1016/0885-2006(95)90005-5
Wittgenstein, L. (1953/2001). *Philosophical Investigations* (G. E. M. Anscombe, Trans.). NY: Blackwell.
Wood, L., Kendal, R., & Flynn, E. (2013). Whom do children copy? model-based biases in social learning. *Developmental Review*, *33*, 341–356. doi: 10/f5hzms
Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, *69*(1), 1–34. doi: 10.1016/s0010-0277(98)00058-4
Woolley, J. (1995). The fictional mind: Young children's understanding of imagination, pretense, and dreams. *Developmental Review*, *15*, 172–172. doi: 10/c9q3qg
Woolley, J., & Ghossainy, E. (2013). Revisiting the fantasy-reality distinction: children as naïve skeptics. *Child Development*, *84*(5), 1496–510.
Wu, C. M., Vélez, N., & Cushman, F. A. (2022). Representational exchange in human social learning: Balancing efficiency and flexibility. In I. C. Dezza, E. Schulz, & C. M. Wu (Eds.), *The Drive for Knowledge: The Science of Human Information Seeking* (p. 169). Cambridge, UK: Cambridge University Press.
Wu, Y., & Gweon, H. (2019). Surprisingly unsurprising! infants' looking time to improbable events is modulated by others' expressions of surprise. *PsyArXiv*. doi: https://doi.org/10.31234/osf.io/8whuv
Wu, Y., & Gweon, H. (2022). Infants Use Others' Surprise as Vicarious Prediction Error. *PsyArxiv*. doi: 10.31234/osf.io/8whuv
Wu, Y., & Schulz, L. E. (2018). Inferring beliefs and desires from emotional reactions to anticipated and observed events. *Child Development*, *89*(2), 649–662. doi: 10.1111/cdev.12759
Wynn, K. (1992). Addition and subtraction by human infants. *Nature*, *358*(6389), 749–750. doi: 10.1038/358749a0
Xu, F., & Garcia, V. (2008). Intuitive statistics by 8-month-old infants. *Proceedings of the National Academy of Sciences*, *105*(13), 5012–5015. doi: 10.1073/pnas.0704450105
Xu, F., & Tenenbaum, J. (2007). Word learning as Bayesian inference. *Psychological Review*, *114*, 245–272. doi: 10/b57qt9
Yarbus, A. L. (1967). *Eye movements and vision.* Plenum Press.
Yildirim, I., Saeed, B., Bennett-Pierre, G., Gerstenberg, T., Tenenbaum, J., & Gweon,

H. (2019). Explaining intuitive difficulty judgments by modeling physical effort and risk..

Yoon, E. J., Tessler, M. H., Goodman, N. D., & Frank, M. C. (2020). Polite speech emerges from competing social goals. *Open Mind*, *4*, 71–87. doi: {10.1162/opmi_a_00035}

Yu, Y., Landrum, A. R., Bonawitz, E., & Shafto, P. (2018). Questioning supports effective transmission of knowledge and increased exploratory learning in pre-kindergarten children. *Developmental Science*, *21*(6), e12696. doi: 10.1111/desc.12696

Yuhas, A. (2021). Over 6 Years and 211 Spots, a British Man Conquers a Parking Lot. *New York Times*.

Zelazo, P. D. (2006). The dimensional change card sort (DCCS): a method of assessing executive function in children. *Nature Protocols*, *1*, 297–301. doi: 10.1038/nprot.2006.46

Zhai, S., Cheng, S., Moskowitz, N., Shen, M., & Gao, T. (2022). The development of commitment: Attention for intention. In *Proceedings of the 44th annual meeting of the cognitive science society.*

Zimbardo, P. G., & Montgomery, K. C. (1957). Effects of "free-environment" rearing upon exploratory behavior. *Psychological Reports*, *3*(3), 589–594. doi: 10.2466/pr0.1957.3.3.589

Zimen, E. (1982). A wolf pack sociogram. In *Wolves of the world: perspectives of behavior, ecology and conservation* (pp. 282–322). Noyes.

Zosh, J. M., Hirsh-Pasek, K., Hopkins, E. J., Jensen, H., Liu, C., Neale, D., ... Whitebread, D. (2018). Accessing the inaccessible: Redefining play as a spectrum. *Frontiers in Psychology*, *9*. doi: 10.3389/fpsyg.2018.01124

Zuccarini, M., Sansavini, A., Iverson, J. M., Savini, S., Guarini, A., Alessandroni, R., ... Aureli, T. (2016). Object engagement and manipulation in extremely preterm and full term infants at 6 months of age. *Research in Developmental Disabilities*, *55*, 173–184. doi: 10.1016/j.ridd.2016.04.001