Tracking changes in exploratory behaviour across development

BIRKBECK, UNIVERSITY OF LONDON

PHD THESIS

Tracking changes in exploratory behaviour across development

Pinelopi-Panagiota Bounia-Mastrogianni 2024

A thesis submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy

School of Psychological Sciences

Birkbeck
UNIVERSITY OF LONDON

Declaration of Authorship

I hereby declare that the work presented in this thesis is my own. All work and materials that are drawn from others are always clearly attributed.

Pinelopi Bounia-Mastrogianni

Aknowledgments

I would first like to express my deep gratitude to Denis Mareschal and Rick Cooper for their invaluable guidance and support during this PhD.

I am particularly thankful to Denis, my primary supervisor, for sharing his broad knowledge and his passion for child development with me and for guiding me with kindness, showing me trust and encouraging me to develop and test my own ideas. His consistent motivation, energy and patience during these years have been crucial for completing my PhD. I am also extremely grateful for the CBCD intellectual environment, which he has played a central role in creating, and has made my PhD a unique, scientifically stimulating experience.
I am also grateful to Rick, my secondary supervisor, for his encouragement, his insightful notes and alternative perspectives on my research and his overall contribution to this PhD.

Thank you, Leslie, Berta, Tamsin, Sarela and Paola, for making the lab run so smoothly on all levels. I am very grateful for the enthusiasm with which families took part in my studies. A special thanks to Ruby and Dian for their invaluable help in data collection, coding and analysis.

I was more than lucky to find wonderful friends among my CBCD lab mates, who made this journey much more enjoyable: thank you Louisa, Lisanne, Eirini, Brittney, Giulia, Olatz and Cécile. Special thanks to Aude for having my back from day one and guiding me on the methodological aspects of my PhD, but most importantly for being a person to share concerns and joy.

Being part of the MOTION network made my PhD even more exciting. Thank you Sabine Hunnius, Kwan Lee and all my fellow ESRs for making my late onboarding smooth and the MOTION experience very rewarding even in the strange pandemic times. It was a delight to do teamwork as part of the network and special thanks to Iara, Joanna, and especially Tommaso and Francesco for the stimulating research collaboration, the inspiring conversations on science and the overall support and friendship.

I am deeply grateful to the special people who made London a place to call home these past five years, and my decision to relocate for the PhD a fullest life experience. Marina, for actually turning a house into a home with your warmth and cheerfulness. Stelio, for giving me perspective and always grounding me with your deep commitment to community. Michael, for the support and validation in the less fun moments. Athina, for transmitting creativity and energy during the last difficult writing months. Christina, for being a mentor and a friend and the most giving person I have met. Finally, Marianna, for being an endless source of stability and the closest to family I found in an unfamiliar place.

Across many borders, a shoutout to my people, who made distance irrelevant in order to be present at every turn of the PhD journey, and most importantly my life journey throughout these years. Saranti, thank you for being there during the past 14 years and for the certainty that you will be in the years to come. Theodora and Lydia, for the irreplaceable tenderness and care. Antoni and Vassili, for consistently reconnecting me, each in a unique way, with my deeper self, worth and values. Nadia, for always showing up and for always being a woman to look up to. And Vassili, for bringing me out of my comfort zone and closer to my desires when I needed it the most, and for keeping it so ever since.

Closing, I want to deeply thank my family, who backed my effort at all stages – my dad, for his unconditional belief in my abilities, but especially my two women: my sister Zoe, for her radiating energy, encouragement, and for making me look to the future with excitement, and my mum, for being strong and bold and for always looking for a way to be there, at all times.

Abstract

There is vast amount of empirical evidence that childhood is a period of increased curiosity for the world and of broad exploration. This exploratory behaviour is manifested both as directed information to resolve uncertainty in the environment, as well as novelty seeking. It has been proposed that both exploratory tendencies gradually become narrower as adulthood is reached, giving space to more exploitative, goal-directed behaviour. However, findings to date have been contradictory, and the exact balance between exploration and exploitation, as well as between exploratory behaviours across development have yet to be clarified. A substantial part of this thesis focuses on the real-time conflict between these options when people are interacting with the world, and how this might change with improving cognitive control across development. We approach this question first by employing hand kinematics analyses in a decision-making task. Here, the analyses of kinematic parameters were found to capture meaningful online decision-making processes in children. A second part of this thesis focuses on information in the physical world, and how the available amount of information might influence object manipulation and exploratory behaviour, specifically when object complexity varies. We find that children prefer novel stimulation more as compared to the other age groups, especially when additional cognitive load is enforced by the decision context or by individual level of executive functioning skills. Finally, we also found that object complexity differentially affects preschoolers' interest and explicit preferences, especially in the visual as compared to the haptic domain. Object complexity also significantly affected young children's still-developing object fitting skills, leading them to use their hands as attentional anchors in the environment. In summary, this thesis shows that humans track informational changes with their perception and action from a very early age (3 years old) and assign value to this information in different ways as they grow older, based on their level of cognitive control abilities, their individual preferences and their contextual or long-term goals.

CONTENTS

List of figures

List of tables

Chapter 1

Extrinsic reward vs. information as motive for action across development: Theoretical and methodological foundations

1.1. Introduction

Human desire for information needs no scientific proof: from activities explicitly related to knowledge, such as reading the news, looking up an unknown word and visiting museums, to more implicit ones, such as playing games, solving crosswords and enjoying forms of art, humans crave for and enjoy acquiring new knowledge about themselves and the world. However, humans avoid information equally often; for example, when they expect bad personal or medical news or when a skill is too difficult for them to master. Every single moment, humans predict and calculate the effects of information on their cognition, their mood and their possibility of gaining rewards. This is then used to guide their actions towards or away from knowledge.

People seek information when they expect to resolve uncertainty and learn. This learning can be associated with high instrumental utility – i.e., knowledge which might increase their possibility to acquire external rewards in the future (such as in explore-exploit tasks, e.g., Wilson, Geana, White, Ludvig, & Cohen, 2014). Or, it can be valuable in and of its own (e.g, Chater & Loewenstein, 2016; Cogliati Dezza, Yu, Cleeremans, & Alexander, 2017; Crupi, Nelson, Meder, Cevolani, & Tentori, 2018; Golman & Loewenstein, 2018; Gottlieb, Oudeyer, Lopes, & Baranes, 2013; Kidd & Hayden, 2015; Kobayashi, Ravaioli, Baranes, Woodford, & Gottlieb, 2019; Oudeyer, Lopes, Kidd, & Gottlieb, 2016; Schwartenbeck et al., 2019; van Lieshout, de Lange, & Cools, 2021). This vast literature on intrinsic motivation, curiosity and active learning will be extensively discussed next in this chapter.

Moreover, the decision to acquire information can significantly impact humans' affect. It has been widely shown that people are more likely to look for positive than negative news and their expectations about how they will be feeling after they learn significantly influences their decision to search out information (Charpentier, Bromberg-Martin, & Sharot, 2018; Cogliati-Dezza, Maher & Sharot, 2022; Karlsson, Loewenstein, & Seppi, 2009; van Lieshout, de Lange, & Cools, 2020; van Lieshout, Traast, de Lange, & Cools, 2021). Taking all these motives into account, decisions on whether to seek information are always based on the context and the relative – and cumulative – utility of available options.

In this thesis, we are interested in the relationship between non-instrumental and instrumental behaviour, as well as the ways they motivate action across development. Next in this chapter, we will discuss key topics on information seeking, how non-instrumental and instrumental behaviour interact and how maturation of cognitive control might influence the balance between the two. Finally, at the end of this chapter, we will discuss our methodological approach.

1.2. What motivates information – seeking?

1.2.1. Novelty vs. complexity-based exploration

The theories about the motives for information-seeking can be grouped in two big categories regarding the stimuli or situation characteristics that excite more curiosity and exploratory behaviour: (i) novelty-based theories and (ii) complexity-based theories. Novelty theories propose that people explore more what they know less about or what they are less confident about (Dubey & Griffiths, 2020). Berlyne (1950) suggested that exploring novel stimuli is intrinsically rewarding and described novelty as a driving force that motivates exploration and diminishes with exposure. Supporting this theory, he showed that rats tend to explore more the stimuli that are more novel to them (1966). Similar findings have been reported in young children and infants; e.g., Smock and Holt (1962) showed that children played more with novel than familiar toys and Fantz (1964) observed that infants direct their attention to stimuli that are maximally novel.

Recent advances in reinforcement learning algorithms have also shown that novelty-based strategies are efficient in exploration tasks (Lehman & Stanley, 2011; Tang et al., 2017; also Twomey & Westermann, 2018 for similar findings using a neural network approach). This is specifically the case when agents have to explore in very unpredictable environments, where new information is highly useful and leads to increased learning (Brändle et al., 2020). However, novelty-based exploration can prove suboptimal in other contexts, as the more novel stimulus is not necessarily the one with the largest subjective utility. For example, exploring based on novelty can lead in learning traps; i.e., situations where no learning can be achieved (Gottlieb et al., 2013). Furthermore, novelty-based exploration cannot explain why people often avoid new stimuli or highly uncertain situations (Kidd et al., 2012; Loewenstein et al., 2001).

Alternatively, it has been proposed that people engage in more information-seeking when stimuli or situations are neither too simple nor too complex. Berlyne (1960) proposed that people feel more curious and explore more when they face an intermediate amount of uncertainty, complexity or incongruity (what he called "collative variables"). In a similar approach, Loewenstein (1994), in his information-gap hypothesis, suggested that people explore more when they become aware of a gap in their knowledge. This approach implies that people become more curious when they already have some amount of knowledge about a stimulus or a topic, but this curiosity diminishes when they know very little or too much about it. This hypothesis has been supported by various studies, showing that curiosity is an inverted U-shaped function of confidence, and that people show the highest curiosity for topics that they are moderately confident about (also termed the *Goldilocks* principle; Kidd et al., 2012). This inverted U-shaped function is also suggested by complexity theory (Dember & Earl, 1957; Kidd & Hayden, 2015), according to which the level of curiosity depends on how well a system can assimilate new information. Thus, the best learning target for a system is neither overly simple (already encoded into memory), nor "too disparate from existing representations already encoded into memory" (Kidd et al., 2012). This idea is also similar to the zone of proximal development, put forward by Vygotsky, according to which students learn optimally when an instructor provides them with knowledge which is neither already learned, nor too far from their accomplished level (Metcalfe et al., 2020; Vygotsky, 1978).

How confident people are in these approaches depends on their prior knowledge of a topic, as well as their expectations about their ability to learn in specific contexts. Such an approach presumes the existence of an evaluation/appraisal mechanism. Thus, recent theories have explicitly suggested that confidence is calculated when encountering an information gap. For example, Grüber and Ranganath (2019) in their Prediction, Appraisal, Curiosity and Exploration (PACE) framework, suggest that people calculate the probability of their information gaps being resolved by incoming information and that this evaluation can lead either to approach ('curiosity') or avoidance ('anxiety'). Similarly, Silvia (2005) proposes the calculation of a 'coping potential' in the face of new information, an appraisal process which can direct people towards or away of new information, based on whether learning can be achieved. These frameworks both point towards curiosity being better understood as a metacognitive process (Goupil & Proust, 2023).

Experimental findings have corroborated these proposals. It has been shown that adults usually explore more when they have low confidence (Desender, Boldt, & Yeung, 2018). For example, Desender, Murphy, Boldt, Verguts, and Yeung (2019) trained an algorithm to classify high versus low confidence responses from electroencephalographic data. Their algorithm could accurately predict whether participants would seek information or not. Similarly, Kang et. al (2009) found that adults were more curious when they reported medium amounts of confidence about an answer in trivia questions. Baranes et. al., (2014) also showed that, when given the choice to organise their practice, people first explore easier tasks and then gradually progress to harder tasks, after having accumulated experience. Children have also been shown to engage in information-seeking based on their confidence levels (Coughlin, Hembacher, Lyons, & Ghetti, 2014; Goupil & Kouider, 2019; Lapidow, Killeen, & Walker, 2022). For instance, in a perceptual identification task, 3- to 5-year-old children asked for additional information – instead of responding by themselves – more frequently in conditions in which they also reported low confidence (Coughlin et al., 2014). In the language domain, 4-year-olds' were more curious about word meanings when their confidence was lower (Jimenez, 2018; Jimenez, Sun, & Saylor, 2018) and even 20-month-olds asked for help when they forgot the location of a toy (Goupil, Romand-Monnier, & Kouider, 2016).

All of the aforementioned examples suggest that people approach information as a function of their confidence, consistently monitoring their knowledge state. As mentioned above, apart from their prior knowledge, people also simultaneously monitor the possibility of gaining information from what is presented to them in the environment. For example, infants have been shown to direct their attention to stimuli based on how complex, predictable or informative they are. Kidd et al. (2012)

showed that infants preferably attend to image sequences of medium unpredictability, which are neither too easy nor too difficult to process. The same finding was observed with auditory stimuli (Kidd, Piantadosi, & Aslin, 2014) and with macaque monkeys (Wu, Blanchard, Meschke, Aslin, Hayden, & Kidd, C., 2022). Infants have also been shown to track the informativity of stimuli and orient their attention to stimuli they can learn from. Addyman & Mareschal (2013) showed that 5-month-olds disengage from visual sequences when the information presented become redundant. Similarly, it was recently shown that 8-month-olds prefer visual stimuli which are associated with larger information gain (Poli et al., 2020).

Some computational approaches have attempted to quantify the notion of 'intermediate' complexity, using information-theoretic measures such as surprise, uncertainty, and information gain (e.g., Kidd et al., 2012, calculated the negative log probability/surprise of expected events). One recent influential approach originating from the field of developmental robotics is the Learning Progress Theory (LPT) which proposes that the brain, as a predictive machine, "is intrinsically motivated to pursue activities in which predictions are improving" (Luciw, Kompella, Kazerounian, & Schmidhuber, 2013, Oudeyer et al., 2016; Figure 1.1.). According to LPT, humans – and artificial agents – can preferentially choose tasks that are learnable, by constantly calculating their prediction errors and directing their exploration based on the minimisation of these errors. This approach often generates different predictions of exploratory behaviours as compared to the predictions derived from novelty-based theories, since the more novel stimulus is not always the one which contributes more to learning. Indeed, recent findings in adults and infants confirm this distinction (Poli, Meyer, Mars, & Hunnius, 2022; Ten, Kaushik, Oudeyer, & Gottlieb, 2021). A related Bayesian approach (Friston et al., 2017; Schwartenbeck et al., 2019) proposes that, based on the availability of information in the environment, people should engage in different exploratory strategies: when information is available and reliable, they should direct their exploration either towards events to-be predicted ("hidden state exploration"), or the causal, parametric structure of the model itself ("model parameter exploration"). When information is not reliable, they should engage in random exploration.

Figure 1.1.

The Learning Progress Theory (LPT) conceptualization of curiosity

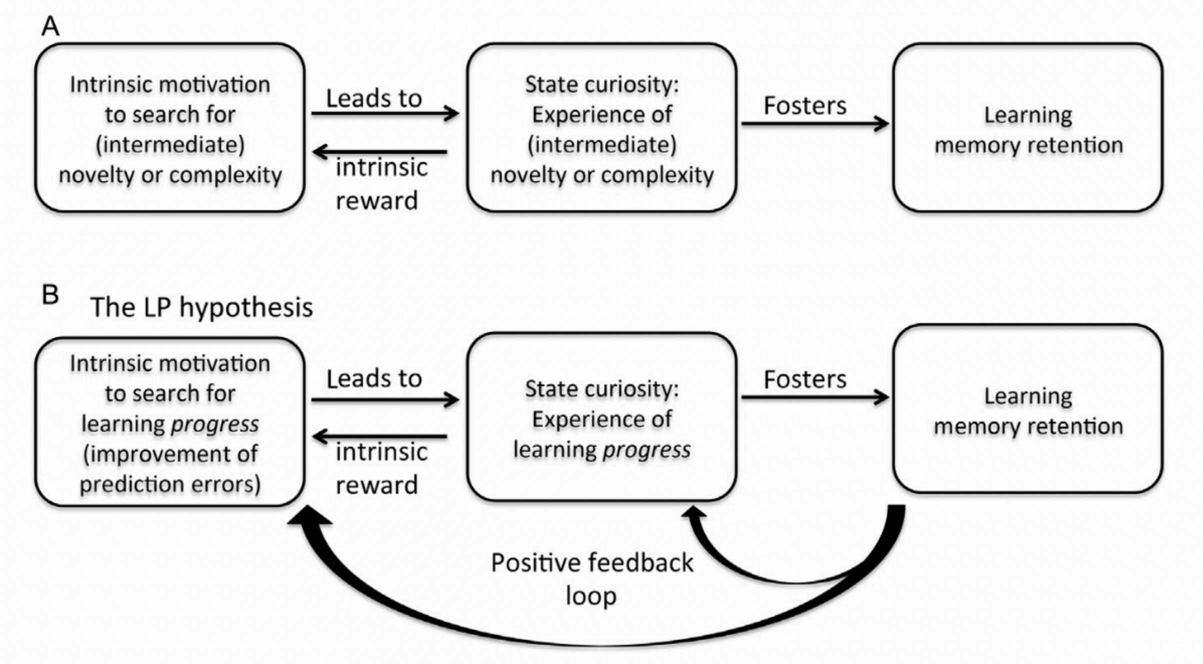

Note. Compared to previous studies of curiosity and learning (A), the learning progress hypothesis suggests that learning progress itself, measured as the improvement of prediction errors, can be intrinsically rewarding, and thus create a positive feedback loop between state curiosity and learning (B). Reprinted from " Intrinsic motivation, curiosity, and learning: Theory and applications in educational technologies" by P.-Y. Oudeyer, J. Gottlieb, and M. Lopes, 2016, *Progress in Brain Research, 229*, p. 266. Copyright 2016 Elsevier B.V.

Complexity-based theories, along with their computational implementations, can accurately explain exploratory behaviour in multiple contexts, but still have limitations. Theoretically, they cannot explain how curiosity is generated and how people will direct their exploration if they have no information about the environment, and thus cannot compare the prediction errors generated by the available stimuli (although, the aforementioned Bayesian approaches propose that random exploration is the chosen strategy in these situations). Related to this, they cannot incorporate situations in which novelty is considered a better predictor of learning progress than prediction errors (i.e., in contexts with very little knowledge about the world).

A recent *rational account* of curiosity and exploration attempts to account for all information-seeking contexts, considering the limited resources that humans have when exploring their environments and their goal to choose the stimulus with the higher knowledge value (Dubey & Griffiths, 2020; Lieder & Griffiths, 2020). According to this view, the volatility of the environment is also taken into account. In stable environments, humans direct their exploration towards areas of medium uncertainty, aiming to minimise their prediction errors – as they expect this error minimisation to prove valuable in future decisions. In contrast, when environments are constantly changing, exploration should favour novelty and surprise. This ensures that humans will choose the option that makes better use of their cognitive resources and will consequently provide higher rewards in the future. Indeed, recent findings in adults and children suggest that they take the cognitive cost of exploration into account when choosing their actions (e.g., Aguirre et al., 2022, showed that toddlers could plan their epistemic actions taking into account whether information would be available in the environment). We will be discussing the role of cognitive control in exploration across development later in this chapter.

In summary, different accounts have described the relative role of informational attributes such as stimulus novelty, subjective uncertainty and incongruity in provoking curiosity and directing exploration. Although much debated, many theorists, as well as direct phenomenological experience, suggest that these attributes are associated with specific emotional states that motivate different goal-directed behaviours. The subjective quality of these emotions has also been extensively discussed and studied, and we will be covering this topic next.

1.2.2. Exploration and affect: how does it feel to be curious?

Although not included in the traditional list of emotions (e.g., Ekman, 1999), curiosity has often been described by recent theorists as an epistemic (Vogl et al., 2020) or a metacognitive feeling (Goupil & Proust, 2023). In reality, curiosity has been approached as a feeling since the early years of its scientific study, with most of the discussion focusing on its valence, while specific categorisations seem to correspond directly with novelty-based vs. complexity-based exploratory motives. Berlyne (1954) described curiosity as an appetitive drive for knowledge, following previous conceptualisations (e.g., Freud, 1915, had described it as a 'thirst for knowledge'). He then went further to differentiate between *specific* curiosity, which arises when people are missing information – similar to the concept of intermediate uncertainty – and *diversive* curiosity, which describes a tendency to look for new tasks/learning goals driven by feelings of boredom (or lack of arousal, e.g., Berlyne, 1967) – a positive feeling which can possibly relate to exploring based on novelty. Two subsequent theories expanded on these directions: Loewenstein (1994; Golman & Loewenstein, 2018) described curiosity as the aversive, unpleasant feeling experienced when people become aware of an information gap that motivates them to fill this gap in – he specifically suggested that its intensity varies as a function of the gap size, with smaller gaps giving rise to feelings such as the 'tip-of-the-tongue'. On the other hand, Spielberger and Starr (1994) in their optimal stimulation model suggested that people explore in order to induce pleasurable curious states to themselves – although approaching dangerous stimuli can also induce anxiety (too much arousal). Building on these theories, Litman and Jimerson (2004) formulated their interest/deprivation(I/D) model of curiosity. Although this model aimed to examine stable traits regarding the ways people approach new information, it can be applied to fleeting emotions in a similar way. According to this model, curiosity can be a positive motivation for learning and engagement (the interest factor) or a negative emotion generated by uncertainty or conflict (the deprivation factor) – both feelings lead to information-seeking but they arise in different contexts. Moreover, different people also exhibit more stable dispositions (trait-like tendencies) towards uncertainty, characterising their usual affective experience (i.e., they might be more or less risky or ambiguity-averse in life in general). The relationship between the exploratory behaviour and these emotions is also quite clear, especially in the case of complexity-based exploration and feelings of deprivation. In the case of novelty-based exploration, the interest factor of the model has been associated with sensation-seeking and openness-to-experience (Kashdan et al., 2018; Litman, 2005; 2008), both factors related to seeking novel experiences, again as more stable individual dispositions[1]. Litman (2005) also draws connections between the I/D model and the wanting/liking dissociation, which refers to the two subcortical neurobiological systems that seem to underlie appetitive motivation and subsequent

[1] The definitions of novelty-seeking do not directly correspond with the ones discussed in the previous chapter section, where it was used as a measure of previous exposure to stimuli. However, the process of directing exploration based on novelty is similar and presupposes looking for stimuli for which there is none or minimal information – the scope of the theories is different.

experiences of pleasure (Berridge, 1999; Berridge & Robinson, 1998). He specifically suggests that the desire for specific information based on deprivation relates to *wanting states* (controlled by the dopaminergic system), while open-ended seeking based on interest is motivated by the *liking motive* (controlled by the opioid system). The two systems always interact; for example, there is always the pleasure of filling-in a gap that might motivate information-seeking but they can also be dissociated (Litman, 2010; Figure 1.2.). The reward-learning framework proposed by Murayama (2020) makes a similar point about learning how to engage in information-seeking.

Figure 1.2

Wanting vs. Liking motives in information-seeking

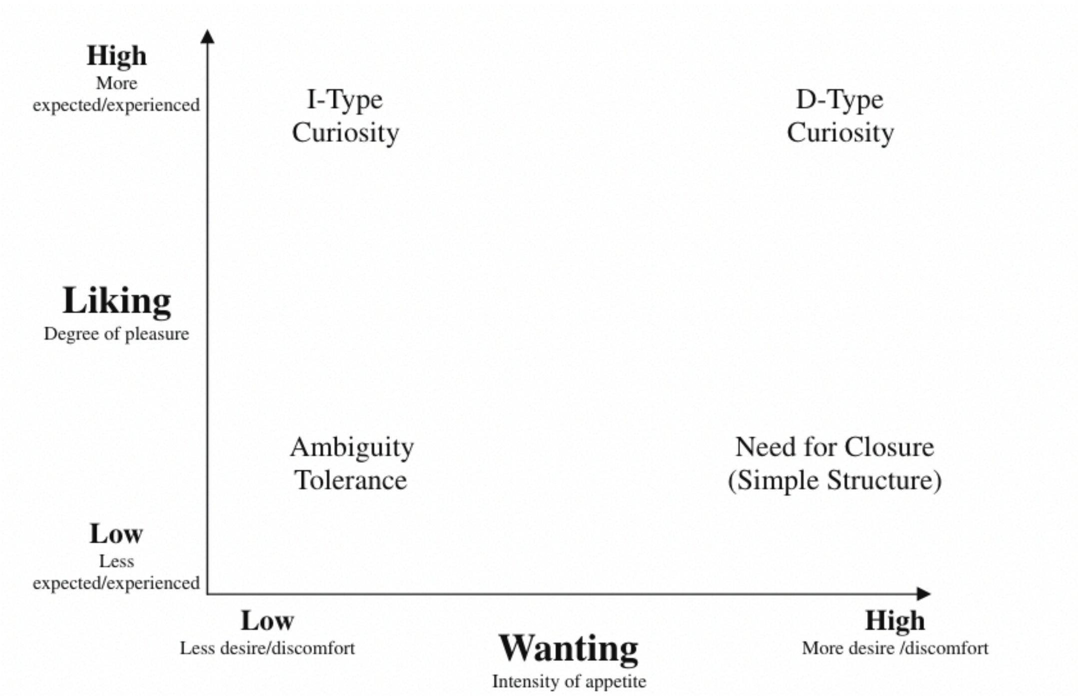

Note. Hypothesized emotional-motivational tendencies relevant to information-seeking associated with relatively high or low levels of wanting and liking. Reprinted from " Relationships between measures of I- and D-type curiosity, ambiguity tolerance, and need for closure: An initial test of the wanting-liking model of information-seeking," by J.A. Litman, 2010, *Personality and Individual Differences, 48*, p. 398. Copyright 2009 Elsevier Ltd.

In summary, the discussion on the emotional aspect of information-seeking is directly related to its relationship with the reward system of the brain. The overlapping mechanisms that support both extrinsic and intrinsic rewards will be discussed in more detail in the relevant experimental chapters.

1.2.3. Instrumental vs. non-instrumental information seeking

So far, we have discussed information-seeking and its phenomenology without referring much to humans' motivation when they seek information. Such motivation can be instrumental (i.e., aim in collecting more information which will lead to the acquisition of extrinsic rewards), or non-instrumental (i.e., acquire knowledge for its own sake). Factors such as novelty and complexity influence humans' information-seeking similarly regardless of their motivation. For example, in explore-exploit tasks (Cohen, McClure, & Yu, 2007; Mehlhorn et al., 2015), where information-seeking is directly related to increasing one's rewards, people usually engage in directed exploration, choosing the option which will optimally resolve their uncertainty about the available rewards; this strategy takes into account both novelty and complexity depending on the context (e.g., see Gershman, 2018, for a discussion of the algorithms which underlie the exploratory strategies in explore-exploit tasks). While novelty-seeking is considered a more efficient strategy when the environment is volatile and the knowledge about the environment is little; as we discussed above (Dubey & Griffiths, 2020), this is still a different strategy than random exploration in explore-exploit paradigms; the latter likely reflects behavioural variability or a total lack of knowledge (or just temporary disregard) of the values of options (Wilson et al., 2021).

However, most differences between instrumental and non-instrumental information-seeking seem to regard the affective part: theorists refer to feelings of curiosity more often when it comes to non-instrumental information-seeking (Gottlieb & Oudeyer, 2018). Non-instrumental information seeking is characterised by the lack of external reinforcers, and is thus considered an intrinsically motivated activity. Intrinsic motivation refers to humans' tendency to seek out information and challenges spontaneously, to independently choose to expand their skills and knowledge without the presence of external rewards (di Domenico & Ryan, 2017). It has been suggested that intrinsically motivated learning serves adaptive functions in organisms, allowing them to survive in changing environments, whereas innate mechanisms take longer to develop and often prove too limited (Baldassare, 2013; Blain & Sharot, 2021).

The advantages of built-in self-motivated learning abilities have been recently shown in artificial agents' adaptivity (e.g., Colas, Fournier, Chetouani, Sigaud, & Oudeyer, 2019; Forestier, Mollard & Oudeyer, 2016). On the phenomenological level, the inherent subjective value of non-instrumental learning in humans proposedly originates from feelings of increasing growth and fitness. Ryan and Deci (2000, 2017; also, di Domenico & Ryan, 2017) with their self-determination theory (SDT) suggest that intrinsic motivation supports two basic psychological needs: the need for autonomy and the need for competence. Based on DeCharms (1968), they define autonomy as a state of volition characterised by a feeling of authenticity and self-directedness in one's actions, as opposed to internal conflicts, pressures, or external coercion, while competence is defined as the experience of effectance, which involves a perception of steadily improving proficiency in tasks that offer an optimal level of challenge and contribute to the development of one's abilities. Similarly, Blain and Sharot (2021) suggest that actions are satisfying in and of themselves when they contribute to increased self-efficacy - a similar concept to the ones proposed by Ryan and Deci which incorporates feelings of agency and competence. Indeed, the experience of agency and competence during learning have been shown to produce positive feelings; they seem to relate to the concept of flow (Csikszentmihalyi, 1990), which refers to a state of performing and learning in an optimal, smooth, automatic way. This positive feeling, when experienced during learning, likely reinforces the information-seeking process, adding value to it in a similar way as external rewards do (Murayama, 2022). This relationship between intrinsic and extrinsic reward, as well as the way it motivates behaviour, will be discussed in the following chapter section.

1.3. Extrinsic vs. intrinsic reward: Similarities and differences

We have already shown above that information can motivate (exploratory) action. It has been suggested that this information seeking motivation resembles the way that other actions are motivated be external reinforcers (i.e., food or money): if an action leads to a rewarding experience, its subjective value increases, and thus so does the likelihood of repeating it. The subjective value of information seems comparable to the value that originates from extrinsic rewards. Many studies have shown that humans and primates treat information as rewarding for its own sake and that they are even willing to sacrifice food or money for information. For example, Charpentier, Bromberg-Martin and Sharot (2018) and Vellani, de Vries, Gaule and Sharot (2020) found that people were willing to pay to get information in advance about the outcomes of a gamble, especially when they expected these outcomes to be positive. Similarly, Bromberg-Martin and Hikosaka (2015) showed that macaque monkeys were willing to sacrifice water to get advance information about winning more water in a comparable task. These studies have also examined the neural mechanisms that underlie both extrinsic and intrinsic reward-based behaviour, and potential overlap has been identified. These (possibly) shared mechanisms might also explain how people learn to assign value to information in the long-term scale.

1.3.1. Shared and separate mechanisms

While explicit behaviour and affect seem to suggest that extrinsic and intrinsic rewards work in a similar way, whether the brain uses the exact same mechanisms and computations is still not fully understood. The fact that humans and recent evolutionary ancestors engage in intrinsically rewarding activity much more often than simpler organisms might imply that these activities involve the neocortex and do not rely on the dopaminergic midbrain system. Another hypothesis is that a combination of systems might be at play, or that the intrinsic rewards relies on the same primary rewards system to guide behaviour based on learning. Several studies point towards an overlap; for example, single-cell recordings in monkeys (Blanchard, Hayden, & Bromberg-Martin, 2015) show that the orbitofrontal cortex encodes both possible primary reward and informativeness, whereas neuroimaging and pharmacological studies in humans (Charpentier et al., 2018; Lau et al., 2020; Vellani et al., 2020) have identified that the same systems and neuromodulators are involved in signaling information and primary rewards – they specifically found that information prediction errors (i.e., errors in predicting how informative a stimulus is) are encoded by the ventral striatum, such as primary reward prediction errors. Furthermore, even without direct experimental comparisons between primary rewards and information, many studies have consistently reported activation of the brain's mesolimbic dopaminergic circuit when people are curious, looking for information, or learning new things (Tomov et al., 2020). However, it was shown in monkeys that even within the OFC, distinct neurons encode either the value of primary rewards or of information, suggesting there might be variations in the cell level (Blanchard et al., 2015). Also, while valuation of primary rewards and information might largely overlap, the integration of these calculated values and the subsequent action/decision-making largely involves areas of the prefrontal and motor cortex (Tomov, Truong, Hundia, & Gershman, 2020). The interplay between these areas is yet to be fully understood, but in general, it seems that intrinsically rewarding information seeking activates at least partly the same brain areas that external rewards do.

This large overlap might explain how people learn to value information as they develop, and subsequently use this value to guide their behaviour. In their recent process account of curiosity and interest, Murayama and colleagues (2019; Murayama, FitzGibbon, & Sakaki, 2019) propose that humans learn the value of information as they continuously engage in information-seeking and resolve their uncertainty, based on a positive feedback loop similar to the one proposed in reinforcement learning: the positive feeling experienced by the resolution of uncertainty is rewarding, thus leading to more information seeking, just as the acquisition of primary rewards and the accompanying pleasure reinforce such seeking behaviours. According to this framework (Figure 1.3.), information holds *incentive salience*, a term used to describe the motivational character of rewards in the reward-learning literature (e.g., FitzGibbon, Lau, & Murayama, 2020). This conceptualisation is consistent with several theories discussed previously; e.g., the LPT, which also suggests that learning itself motivates information-seeking in a reinforcing loop.

Figure 1.3

The incentive salience of information

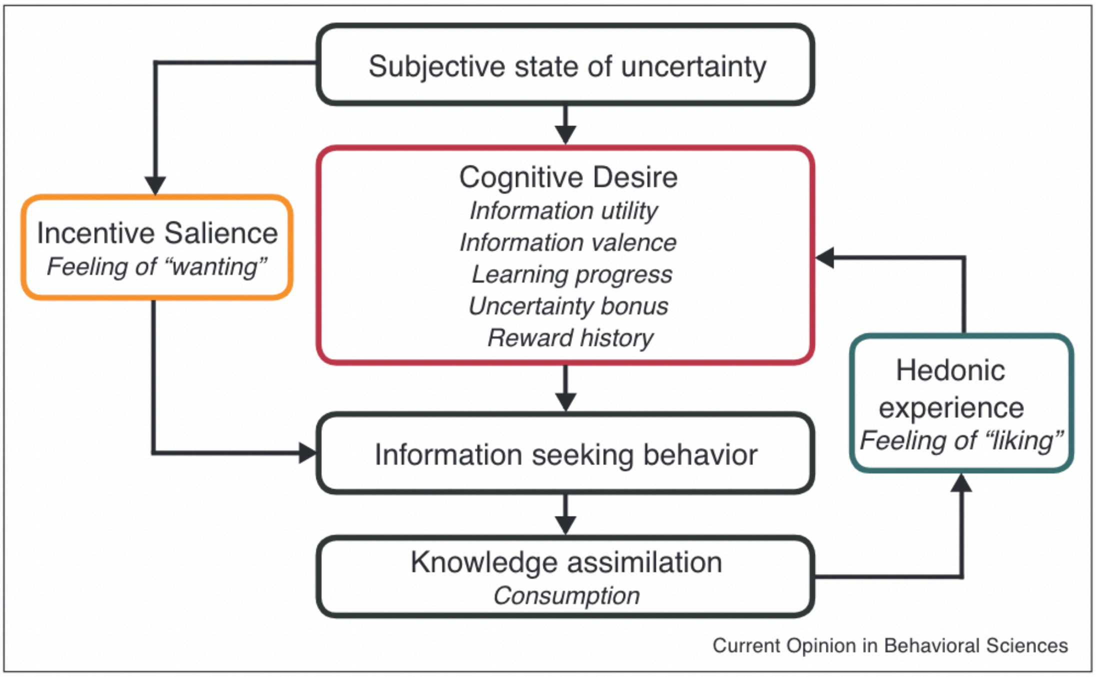

Note. The knowledge acquisition process as reward learning supported by incentive salience. Reprinted from " The seductive lure of curiosity: information as a motivationally salient reward," by L. FitzGibbon, J.K.L. Lau and K. Murayama, 2020, *Current Opinion in Behavioural Sciences, 35*, p. 23. Copyright 2020 The Authors.

In summary, extrinsic and intrinsic rewards seem to share computational and neuronal mechanisms and generate similar motivational states. As we are interested in studying these processes across development, we will next be discussing how they manifest in different ages, and possible hypotheses explaining these differences.

1.4. How does exploratory behaviour change across development?

While extrinsic and informational value motivate action throughout life, common experience and scientific evidence show that humans might differ in the relevant amount and strategies of exploration they engage in at different developmental stages. Children and adolescents are commonly thought to be more exploratory and driven by new experiences than adults. However, the specific changes in behaviour and possibly brain functions are still under examination.

Humans can track informativity in the environment and direct their attention to it from infancy (Addyman & Mareschal, 2011; Kidd et al., 2015; Stahl & Feigenson, 2015; Poli et al., 2020; Wu, Gopnik, Richardson, & Kirkham, 2011), while toddlers and preschoolers recognise uncertainty and learning possibilities (e.g., Aguirre et al., 2022; Ruggeri, Pelz, Gopnik, & Schulz, 2021) and engage in exploratory action to de-confound variables (Schulz & Bonawitz, 2007), to examine belief violations (Bonawitz, van Schijndel, Friel, & Schulz, 2012) and test their intuitive theories and hypotheses (Cook, Goodman, & Schulz, 2011). Despite this general sensitivity to informativity, it has been suggested that children's exploration differs from adults in breadth (i.e., children engage in broader exploration; Gopnik, 2020), and possibly to differential sensitivity to parameters of information: infants and young children are possibly affected differently by novelty compared to uncertainty in the environment (Goupil & Proust, 2023; Nussenbaum et al., 2022).

The specific trade-off between exploration and exploitation has been extensively studied with traditional bandit tasks (Sutton & Barto, 1998), which can capture human behaviour when faced with the dilemma between extrinsic rewards and information[2]. Such tasks can also particularly shed light on the shift in exploratory strategies used at different ages. In many such examples, children have been found to explore more broadly first, and settle later for the more rewarding option, a strategy which often protects them from 'learning traps' and allows them to uncover more complex rules in unknown environments. For example, Liquin and Gopnik (2022) compared preschoolers' and adults' performance in an approach-avoid task and found that young children explore for longer and are thus more accurate in understanding the complex structure of the presented environment. Similarly, children as young as 3 years old choose equally often the bad and good options in bandits throughout the tasks (Blanco & Sloutsky, 2021; Sumner, Steyvers, & Sarnecka, 2019). Equally, 5- to 12-year-olds are better at noticing changes in the structure of a multi-armed bandit task than adults (Sumner et al., n.d.). It has recently been shown that this advantage might result from more distributed

[2] Information-seeking in explore-exploit tasks is almost exclusively instrumental – tied to gaining extrinsic rewards, as we already discussed in 1.3. However, we consider the findings we present here relevant to the developmental changes in exploration in general.

attention during the learning phase in a category-learning task (Blanco, Turner, & Sloutsky, 2023), but this might also reflect different information processing during exploration of the environment more generally.

Furthermore, evidence suggest that children differ from adolescents and adults in the relevant amount of directed and random exploration they engage in when searching a new environment (Gopnik, 2017; Wilson, Bonawitz, Costa, & Ebitz, 2021). Directed (or systematic) exploration refers to exploratory behaviours directed towards informative parts of the environment. They are accompanied by an information or learning progress 'bonus' that makes the exploratory choice more valuable. In contrast, random exploration reflects behavioural variability (i.e., noise-in-choice behaviour) and it is more common when there is no prior knowledge about the values of options (extrinsic-reward or information-wise). Studies which have used variations of bandit tasks show that children engage in more random exploration than adolescents and adults. For example, Wu, Ruggeri and Meder (2019) compared children in middle and late childhood to adults in a spatially correlated multiarmed-bandit task. Their findings showed that children engaged in more directed exploration and generalised less than adults, but the groups did not differ in terms of random exploration. The same research group (Meder, Wu, Schulz, & Ruggeri, 2021) used a similar task to disentangle children's exploration strategies at 4 and 9 years of age, showing that random exploration decreased as children grew older. They also found evidence of directed exploration in the youngest group. Somerville et al., (2017) investigated 12 to 28 year-olds, using armed bandit tasks to manipulate the time horizon of the informational utility. In their study, adolescents appeared more influenced by immediate rewards. Moreover, adolescents who engaged in more random exploration also scored higher in a risk-taking scale.

The aforementioned tasks differentiate between exploration based on uncertainty/information value and exploration without knowledge of values. However, exploration might also differ based on the type of informational attribute which might be more relevant for learning in different age groups or different contexts. For example, novelty-seeking might underlie both exploration based on directed information seeking as well as some behaviours typically considered as random exploration.

Recent studies have shown that stimulus novelty has a separable influence on exploratory behaviours compared to uncertainty or expected learning progress in adults (Cockburn et al., 2021; Poli et al., 2022). Furthermore, Nussenbaum et al. (2022) compared participants ranging from 8 to 27 years of age in an exploration task that separately manipulated stimulus novelty and reward uncertainty. Interestingly they found that, while all age groups were influenced by novelty in their exploration, children showed no uncertainty aversion and explored the uncertain options more than the adults. This lack of aversion in early childhood has been documented before (Li, Roberts, Huettel, & Brannon, 2017; Rosenbaum & Hartley, 2018) and might explain children's exploratory tendencies. Moreover, Blanco and Sloutsky (2021) document that children visit all the options in an armed bandit task using a time-dependent strategy (i.e., keeping track of the options they have not visited recently and using this information to guide their exploration).

Despite still being under investigation, the general tendency of children towards broader exploration, which possibly still undergoes further change during adolescence, has been related to the protracted human development and reliance on learning. Specifically, Gopnik et al. (2017; Gopnik, 2020) has proposed human's unique developmental life history (i.e., a very long childhood, many carers and the long reliance on them, increased brain plasticity and sensitive periods for a long time) might underlie the developmental changes observed in the trade-off between exploration and exploitation. Gopnik and her colleagues suggest that the lack of mature executive functions, which results in minimised capacities for focused attention, planning and delayed gratification, also allows children to distribute their attention more broadly during learning (e.g., Blanco et al., 2023, as we discussed earlier; Plebanek & Sloutsky, 2017), generate and test broader hypotheses (Lucas, Bridgers, Griffiths, & Gopnik, 2014) and engage in more creative, 'divergent' thinking (Thompson-Schill et al., 2009). Indeed, children have been shown the ability to imagine more new uses for a tool (German & Defeyter, 2000).

In summary, there is compelling evidence to suggest that children explore more and differently than adults. How children differ from adolescents is still unresolved. The exact differences in the role of informational attributes across development, though, is only now starting to be investigated. Moreover, the role of developing cognitive control and the possible individual role of separate executive functions (e.g., inhibition) might play is still, to our knowledge, not understood. We aim to contribute towards this direction with our studies.

1.5. Methods

1.5.1. Paradigms to study curiosity and exploratory strategies across development

We have already discussed various tasks used by researchers to measure curiosity and exploratory behaviour, focusing mostly on situational manifestations of such behaviour. As curiosity and exploration have always been of relevance in education and psychology, there was also interest in their measurement as a stable trait. Jirout and Klahr (2012) review many of these measurements, for example the different questionnaires used in the past decades to measure trait curiosity (e.g., the Ontario Test of Intrinsic Motivation (OTIM) by Day (1971), or the State-Trait Curiosity Inventory (STCI) by Spielberger et al. (1980)), as well as behavioural measures such as the amount of spontaneous exploration of objects (e.g., Minuchin, 1971; Smock & Holt, 1962) and preference for complexity (e.g., Henderson & Moore, 1980). Even though they were used to investigate (proposedly stable) individual differences in curiosity and possible relationships with academic achievement, these behavioural tasks are similar to the ones used to study exploration as a spontaneous behaviour – and possible factors that affect it – in more naturalistic settings (e.g., Schulz and Bonawitz, 2007; Taffoni et al., 2014 and others we have already discussed above). At the same time, extensive research has used variations of bandit tasks to study the balance between exploration and exploitation, and developmental differences in this balance.

Most of these tasks, while very diverse, show important commonalities if we attempt to analyze the underlying cognitive processes that take place. When participants have to choose between interacting with different objects or features (in naturalistic tasks) or between exploiting and exploring options in bandits, they essentially engage in value-based decision making. The steps humans follow when deciding based on subjective value (i.e., the representation of the decision space, the calculation of values, the value comparison and accumulation of evidence; e.g., Tajima, Drugowitsch & Pouget, 2016) as well as the factors which affect the decision process (e.g., time limit; Milosavljevic, Malmaud, Huth, Koch, & Rangel, 2010) can similarly be taken into account when analyzing humans' behaviour in information-seeking contexts, but such approach has not been employed so far, at least extensively. Such an approach could help identify how different rewarding options influence action, and possibly also help dissociate the neural correlates which underlie extrinsic and intrinsic reward – for example, it could be the case that people are taking the two types of rewards into consideration on different steps in their decision process. Furthermore, recent experimental findings have suggested that cognitive control plays a significant role in decision-making. Frömer, Wolf & Shenhav (2019) showed that the specific task goals influence participants in dissociable ways compared to how rewarding they consider the options. This interesting observation might be of relevance to our developmental perspective, as the ability to follow goals or prioritize between them is still under development up to adolescence (Somerville, & Casey, 2010). Taking these into consideration, we decided to design a series of tasks where we can track the decision process when participants choose between options with different value (extrinsic vs. different informational rewards) and use mouse and finger tracking to reveal the underlying steps and processes. More details on this methodology will be discussed in the following subchapter.

1.5.2. Decision-making and hand kinematics

Over the past decades, hand-movement tracking methods (3D motion-capture, touchscreen finger-tracking, mouse tracking), have been increasingly utilized in psychological sciences, their spread especially facilitated by recently developed open-source software, particularly in finger- and mouse-tracking (for example *MouseTracker* by Freeman and Ambady, 2010; or *mousetrap* by Kieslich and Henninger, 2017). Experimental findings from various cognitive tasks have provided evidence against a serial, feed-forward, and stage-based view of mental processing, and supported a more dynamic view of the mind in which "processes across perception, cognition, and action often – though, not necessarily always – unfold in a parallel, interactive, and continuous manner" (Cisek & Kalaska, 2010; Lakoff & Johnson, 1999; Smith & Gasser, 2005; Spencer et al., 2009; Spivey, 2007). The use of tracking methods has been particularly fruitful in language processing

(Dale & Duran, 2011; Farmer, Cargill, Hindy, Dale, & Spivey, 2007; Spivey, Grosjean, & Knoblich, 2005; Tomlinson, Bailey, & Bott, 2013), numerical cognition (Dotan & Dehaene, 2013; Faulkenberry, Montgomery, & Tennes, 2015; Marghetis, Núñez, & Bergen, 2014; Song & Nakayama, 2008), reasoning (Travers, Rolison, & Feeney, 2016), and social cognition (Duran, Dale, Kello, Street, & Richardson, 2013; Freeman & Ambady, 2009; Freeman, Ma, Han, & Ambady, 2013; Freeman, Pauker, & Sanchez, 2016). Simultaneously, a different line of research focusing on consumers' choices has also extensively used hand tracking methods, as well as more broadly research on value-based decision-making (e.g., Koop & Johnson, 2011, 2013; Lee & Hare, 2022; O'Hora, Carey, Kervick, David Crowley, & Dabrowski, 2016). This latter field's advances are of interest for our research, especially in terms of their analysis methodology, as it does not involve correct and wrong options and responses, but rather the incorporation of values during the decision-making processes.

While hand-movement tracking has been extensively used with adult participants, its employment in developmental research is limited, possibly due to factors related to manual dexterity and accuracy in younger ages, which might render such methods too noisy. For example, two studies on cognitive control used hand-tracking in 3D space to compare children's (5- to 10-year olds) and adults' inhibitory and switching performance (Erb, Moher, Song, & Sobel, 2017a, 2017b) and revealed improvements in cognitive control in older children, reflected in specific movement parameters. Schroer, Cooper and Mareschal (2021) also used 3D hand-tracking with motion capture to investigate preschoolers' ability to plan action sequences, showing different use of the non-reaching hand between children with different planning skills. Another study used mouse-tracking to identify arithmetic difficulties in an online learning environment (de Mooij, Raijmakers, Dumontheil, Kirkham, & van der Maas, 2020), while a series of mouse-tracking studies looked into the influences of food attributes (healthiness, taste) and self-control on children's decision-making (Ha et al, 2016; Pearce, Adise, Roberts, White, Geier, & Keller, 2020). In total, the use of such methods in developmental research is on the rise. However, there is need for methodological clarity in the design and analysis aspects of such approaches.

Regarding the design of mouse-tracking tasks, numerous studies over the past decade have focused on important factors which affect performance and assumptions, such as starting procedures (time-limited or not), target sizes, response type (hover vs. click), stimulus positions (centred vs. edged), and so on (e.g., Hehman et al., 2015; Kieslich et al., 2020; Schoemann et al., 2019). The analyses procedures have also been discussed, especially regarding the specific parameters of movement which are more or less informative. The proposed measures are very diverse, and vary greatly due to the differences in each task design on the spatial level (e.g., where stimuli appear, how many stimuli, whether there is an optimal path etc.), as well as the cognitive processes which proposedly underly the task according to its theoretical assumptions (e.g., conflict due to competing phonological activations, vs. conflict due to competing reward values). A common approach which is applied in various task designs and theoretical interests is the single feature approach , where various geometrical, temporal and entropy measures are calculated from the spatial coordinates and time data, and then some of them are singled out to proposedly reflect more information about the processes of interest (Figure 1.4. from Wirth et al., 2020). This approach is widely used and accepted, however the high degrees of freedom it allows to researchers leads to lack of consensus regarding the measures and often difficulties in reproducibility.

Figure 1.4

Geometrical, temporal and entropy measures calculated from kinematic data

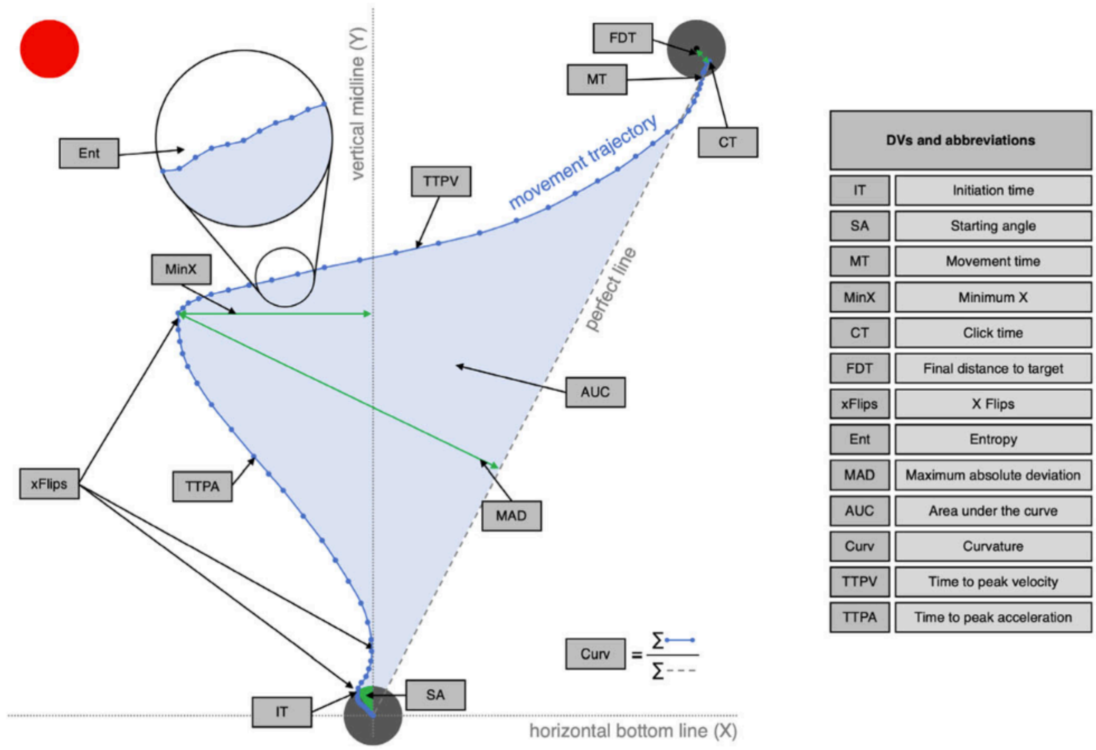

Note. Overview of possible dependent variables when analyzing movement trajectory data. The circle at the bottom represents the starting area; the two on the top left and top right represent the target areas. Reprinted from " Design choices: Empirical recommendations for designing two-dimensional finger-tracking experiments" by R. Wirth, A. Foerster, W. Kunde and R. Pfister, 2020, *Behavior Research Methods, 52*, p. 2399. Copyright 2020 The Authors.

Recently, Maldonado, Dunbar and Chemla (2019) tried a machine learning approach to identify the most informative measures for a linguistic processing and a double negation task. They used kinematics data points (with dimensionality reduced to 13 principal components with a Principal Component Analysis) to train a supervised classifier (Linear Discriminant Analysis, LDA), which could predict two different theoretical explanations of the movement data. Although their approach was successful, it is not easy to apply to other data or designs.

In our tasks, we decided to follow recent studies on value-based decision-making, and calculate commonly used measures, in an attempt to provide comparable results. For example, Koop & Johnson (2013), who focus on preferential choice, use some geometrical measures (MAD, AAD) and an entropy measure (x-flips), while Pearce et al. (2020) analyzed the temporal unfolding of action. We were interested in calculating all types of measures, as preference based on different types of rewards might either unfold gradually and show as deviating trajectories (reflecting an initial choice and continuous suppression of the alternative), or as a two-step incorporation of the values or a dissociation between the value of a goal vs. the value of information, showing up as sudden changes of mind and direction of movements. We will be broadly following the procedure described by Wulff, Kieslich, Henninger, Haslbeck and Schulte-Mecklenbeck (2021), which we will describe in more detail in the relevant chapter.

1.6. The current studies – Aims and summary

The theoretical and experimental advances discussed so far provide evidence for developmental differences in curious exploration. It has been proposed that exploration gradually becomes narrower as adulthood is reached, giving space to more exploitative, goal-directed behaviour. However, findings have been contradictory so far, and the exact balance between exploration and exploitation, as well as between exploratory behaviours across development have yet to be clarified. Substantial part of this thesis focuses on the real-time conflict between explore/exploit options when people are interacting with the world, and how this might change with maturation. We are approaching this by employing hand kinematics analyses in a decision-making task (more in Methods subchapter). A second part of this thesis focuses on information sampling in the physical world, and how the available amount of information might influence exploratory behaviour, specifically manipulating object complexity.

Chapter 2 discusses our first experiment, which aimed to identify how object complexity might relate to motor imagery and self-oriented gesturing, building on the embodied cognition hypothesis that manipulating and imagining objects might lead to similar activation of the motor system. Specifically, 3-5 year-olds manipulated objects of different complexities and their hand movements during mental rotation on a subsequent phase were video-recorded (and partly measured through motion-capture). The experiment identifies some embodied strategies that children use to think about objects, and indeed shows that object complexity linearly increases these strategies. However, no explicit gesturing is documented.

Chapter 3 includes our series of mouse- and finger-tracking experiments (experiments 2-5), in which we aimed to focus on the specific process of planning an exploratory versus an exploitative action and investigate how these decisions unfold real-time. Furthermore, we aimed to examine how its parameters can be predicted by age and individual differences. Our participants chose between different options to interact with, when these options led either to the attainment of a rewarding goal, to missing information or to novel unpredictable stimuli. We looked into how the competition between these options was reflected in the real-time action plans; i.e., at the specific hand kinematics while participants made their choices. Participants completed a computerized decision-making task as we tracked their mouse positions, stimulus-preference tasks with varying levels of uncertainty and complexity and standard executive functions tasks. In different versions of this task, we manipulated specific parameters: In experiment 2, participants had a short exploration horizon, i.e., the acquisition of the extrinsic reward (which required following specific steps) prevented them from exploring the other two options. This was changed in experiment 3, where participants could keep exploring for a longer time. In experiment 4, the experimental stimuli associated with missing information were changed to better capture perceptual uncertainty, and the ones associated with novel stimulation were also replaced with different ones, more unpredictable, preventing any possibility of guessing based on categorization. Eventually, in experiment 5, the extrinsic-reward option was replaced: instead of requiring goal-following, this option was offering rewards on every trial. The experiments overall show differences in participants' preferences, especially when time-limit is imposed, showing that children engage in more novelty-based exploration than older groups. This is also the case for individuals with lower EF skills, although not consistently across experiments. Kinematics data reveal greater conflict between the two exploratory options.

Chapter 4 includes experiment 6, in which we aimed to examine whether preschoolers (4-year-olds) gather information through object exploration in a different manner across the visual and the haptic modality, and whether their interest in exploration can dissociate from their explicit liking. Specifically, children explored sets of 3D-printed objects with three different levels of complexity unimodally through vision or touch, and we measured the exploration time spent on each object and their subjective preferences for the objects. Results showed a linear relationship between visual complexity and exploration time, but no significant relationship for haptic complexity. Explicit preferences were not consistently affected by complexity or modality and did not correlate with exploration time – large individual differences were observed.

Chapter 5 discusses all the experimental findings together, to form a comprehensive broader view of the influence of information on perception, action and subjective valuation across development. In this chapter, we make a point about the need for more nuanced theoretical and experimental manipulations of the concept of novelty. We also emphasize the distinct effects of object complexity on exploration and aesthetic judgements. Finally, we discuss this thesis' limitations and propose new directions for research.

Chapter 2

Action planning and spatial-motor imagery – Effects of object complexity on preschoolers' real and imagined object-fitting

2.1. Introduction

After achieving the important milestone of grasping objects in their immediate environments, infants start performing a variety of prehensile activities that extend beyond this and require planning (e.g., putting toys in containers, piling up blocks to build a tower or putting an item in someone's hand). In order to accomplish such goals, a sophisticated coordination between perception and action is necessary. This is particularly evident when it comes to fitting objects into apertures, a skill widely used as a marker for developed motor problem-solving abilities at the second year of life. In object-fitting, a combination of a wide range of perceptual and motor skills are at play that offer the chance to examine how successfully real-time planning is accomplished.

In order to be able to fit objects into apertures, children have to develop certain perceptual abilities, such as object unity (Johnson & Aslin, 1996), recognizing the objects' different shapes and sizes and the constancy of such attributes (e.g., Soska & Johnson, 2008), the spatial relations between objects (e.g., Rigney & Wang, 2015) and eventually, to understand the shapes and size of apertures (e.g., Adolph, 2000).

However, when it comes to using these perceptual skills to guide action on objects, there are large individual differences between children of the same age in their ability to systematically apply them in relevant tasks. Arguably, this goes back to the evidence of a dissociation between ventral and dorsal visual stream maturation (Johnson, Mareschal, & Csibra, 2008), with the latter having a more protracted development. Many studies have looked into the ways children adjust their reaches and grasps prospectively based on (visually gathered) information on object characteristics, such as size, shape (e.g., symmetric or not) and orientation. Most of these studies suggest that children start making such adjustments from the second part of their first year of life (e.g., Berthier & Carrico, 2010). Although children can configure their hands successfully when they approach an object based on its characteristics, they cannot apply the same skill when they have to insert an object into an aperture until after their 18th month of life. This is often referred to as a ‘coordination of spatial frames of reference’ problem. When acting directly on an object, children use an egocentric frame of reference: they compute their hand/body position in comparison with objects on the environment. However, when holding an object, the hand and the object change position and orientation differently compared to the environment (allocentric frame of reference), adding degrees of freedom and thus, increasing difficulty and cognitive resources’ requirements (Lockman, 2000). Furthermore, the geometric structure of the objects to be fitted into aperatures also influences the difficulty of fitting: asymmetric objects and objects that can only be fitted along a specific axis are more demanding and later to be mastered (Ornkloo & von Hofsten, 2007; Fragaszy, Kuroshima, & Stone, 2015). Despite revealing all these contributing factors in successful object-fitting, these previous studies have not looked into the real-time process of planning actions in order to accomplish the fitting goal, but have rather focused on achievement in different ages.

Taking a more process-focused approach, Jung et al. (2015) had toddlers between 16 and 33 months of age fit a bar of different orientations (vertical, horizontal) into an aperture in front of them and tracked their hand movements with motion capture. Their findings showed that less efficient children under 20 months had a two-stage fitting strategy: first translating the bar (i.e., moving it towards the hole), and then rotating it after they had approached the opening. This suggests that they did not plan ahead to pre-align the object while transporting it. This lack of planning is supported by the fact that they changed the orientation of the bar even in the cases where it already had the right orientation from the beginning (i.e., 50% of the cases). In contrast, more efficient fitters combined translation and rotation while moving the object towards the opening. Crucially, the most efficient solvers kept the object close to the table, avoiding the rotation of the bar in the vertical axis, which was not necessary for the task.

Fragaszy et al. (2015) and Ornkloo and von Hofsten (2007) also focused on how young children (from 14-months up to four years old) processed different object characteristics as they attempted to fit them into holes. Specifically, the more complicated a shape was – thus, having more spatial features to process in order to be aligned successfully with a hole – the more difficult it was for children to pre-align it. The manipulated complexity had specifically to do with symmetry and with the axis of elongation which should be considered to fit the objects, and with their interaction. Specifically, children had more difficulty fitting asymmetric objects (i.e., which could only be fitted with specific orientation along their middle axis), especially when the aperture was matching their short axis (e.g., Figure 2.1a, b; Fragaszy et al., 2015; Ornkloo & von Hofsten, 2007). Due to this added difficulty, children had to explore more options with vision by comparing the objects with the apertures after bringing them above the holes, and with touch by attempting to fit the object in different ways.

All of these studies underline the importance of vision for children to analyse the spatial features of objects and the environment, suggesting that visual inspection plays an important role in formulating the action plan and influencing the online adjustments of hand movements in accordance with a perception-action approach. Both the Jung et al. (2015) study and the Fragaszy et al. (2015) study comment on the direction of visual attention during transport and fitting of the objects, which alternates between the object and the hole. However, none of the studies examines more specifically how this correlates with hand movements in real time.

Figure 2.1

Examples of object complexity manipulation in object-fitting studies, a) example of symmetry stimuli in Fragazcy et al., (2015), b) example of symmetry stimuli with short elongation axes in Ornkloo and vanHofsten (2007).

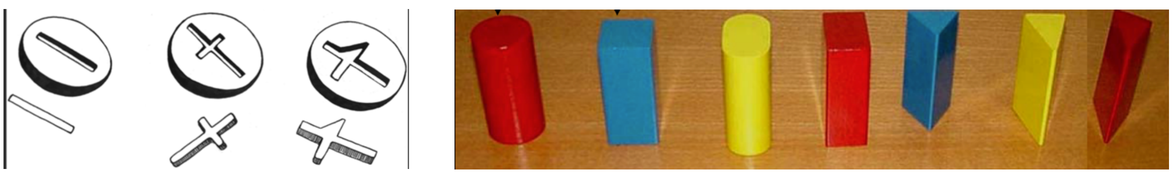

a) b)

This was done by Ossmy, Han, Cheng, Kaplan and Adolph (2020), who used eye tracking to reveal how eye movements unfold during the grasping, transport and fitting of objects, while analysing the hand movements using simultaneous video recordings. This study investigated 3- to 5-year-olds and adults. Interestingly, although they managed to fit the objects in most trials, children showed much less visual attention to their hands during transport of the objects than adults, and only focused on the object and the aperture when they eventually reached the hole. This resulted in delays in the adjustment, which, as the researchers suggest, had to do with delayed sampling of the useful spatial information through vision. Specifically, they suggest that children lack the ability to plan proactively where they should look in order to guide their actions efficiently. Instead, children spend a lot of time looking at irrelevant items in the environment, something that the other studies had not identified.

This inability to direct their vision so as to serve their fitting plan (although they might have already decided correctly which object fits where) could have to do with attentional capture by equally interesting/rewarding stimuli in the environment. Bottom-up attentional capture from salient stimuli has been widely documented in infants' eye-movements (e.g., Gluckman & Johnson, 2013), but it has also been shown to affect hand trajectories in adults (e.g., Song & Nakayama, 2006; Welsh, 2011). This effect on movement has also been reported in the Jung et al. (2015) study in which children move the objects towards the apertures using larger, less efficient routes than the straight ones – but it has not been tested in terms of the exact irrelevant stimuli characteristics in the environment.

Apart from these skills, several investigators have suggested that to pre-align objects with apertures quickly, especially those with a relatively complex spatial structure, children may perform mental rotation (Jung et al., 2015, 2018; Ornkloo & von Hofsten, 2009) and, more specifically, engage in spatial-motor imagery. An interesting window into imagery can be provided by the connection between spatial-motor imagery and hand gestures. From an embodied point of view, gestures are considered to arise from embodied simulations of actions and perceptual events (e.g., the *Gesture as Simulated Action Framework*; Hostetter & Alibali, 2008).

A large number of studies have focused on how different forms of imagery (visual, spatial, motor) give rise to different gestures (Hostetter & Alibali, 2019).

From an embedded/extended cognition point of view, gestures are suggested to provide an external, physical tool which can replace and support parts of the cognitive process (Pouw, De Nooijer, Van Gog, Zwaan, & Paas, 2014). In one of the very few studies that have looked into these processes in children, Wakefield and colleagues found that 4-year-olds benefit more by using gestures than by physical action when they learn how to mentally rotate objects (Wakefield, Foley, Ping, Villarreal, Goldin-Meadow, & Levine, 2019). Furthermore, 4-year-olds spontaneously produce representational gestures (gestures referring to objects/processes) when asked to solve spatio-motoric problems (Boncoddo, Dixon, & Kelley, 2010) and 3-year-olds use private pointing in order to keep the position of a toy in memory (Delgado, Gómez, & Sarriá, 2011). It has also been found that young children who gesture more have better performance in visuospatial tasks, whereas less spontaneous gesturing or inhibition of hand movements can impair performance (O'Neill & Miller, 2013). However, most of these studies have looked into children trying to engage in spatial problem solving and, importantly, have not tried to approach gesturing as an embodied manifestation of the underlying imagery (which must necessarily takes place in children's minds during such tasks). In an attempt to combine the findings of these lines of research, we aimed to (i) understand how preschoolers plan the object-fitting process, (ii) whether they use gestures to aid their mental rotation (specifically spatial imagery) of 3D objects, and (iii) whether these movements are necessary for spatial problem-solving at this age.

We were further interested in exploring the particular parameters of the hand kinematics that can be associated with the real and imagined object manipulation. To this end, 3 to 5 year old children completed a task in which they initially manipulated 3D objects of different complexities and attempted to fit them in an object-fitting toy. They then had to solve the same problem mentally, either with the possibility of using hand gestures, or in a condition in which it was not possible to use hand movements to solve the planning problem. We planned to record their hand movements using Motion Capture and their motor preparation recording the electrical activity of their hand muscles using electromyography (EMG).

Our hypotheses were that: a) the level of object complexity will influence the fitting process, in terms of the needed attempts to fit the object and the fitting time, b) children who plan for a longer time before movement will be more accurate in fitting, c) younger children will plan for a shorter period and will start the transport phase earlier, d) higher spontaneous gesture rates will correlate with the difficulty of the task (shape complexity) and e) successful inhibition of hand movement will significantly decrease the accuracy in the mental rotation task. Regarding the amount of gesturing as a function of age, one possibility is that young children will gesture more, because of lower visual working memory capacity (and, thus, increased need for reliance on proprioceptive/environmental input) and also due to less inhibition of the motor system when they engage in imagery. The particular kinematics of their gestures were expected to correlate with the actual movement kinematics while they manipulated the objects, but we refrain from making more specific predictions regarding these movements, as they might use their hands to mimic either their manipulation or the manipulated object. Finally, we expected to find muscle activity of motor preparation through EMG in the inhibited movement condition (e.g., Addyman, Rocha, Fautrelle, French, Thomas, and Mareschal, 2016 observed such activations in infant EMG data), further supporting the children's tendency to use their hands as support for spatial imagery.

2.2. Methods

2.2.1. Participants

Our participants were 24 children, aged from 3 to 5 year old (Mean: 4.13 years, 9 girls). All participants were neurotypical and had normal or corrected-to-normal vision. The participants were volunteers recruited though the Birkbeck Babylab database.

All children were tested on an object fitting task and a visual-spatial working memory task. Each of these is discussed in turn below.

2.2.2. Object-fitting task

2.2.2.1. Materials

The following objects were used for the experiment: (1) A rectangular house-shaped box with 8 holes on the roof, 4 on each roof side (Figure 2.2a,b). Every hole existed twice in the box, once on each side. The holes are configured in order to fit only one particular object, as a 2D projection of it; (2) The objects, which were different arrangements of 3D rectangles, joined together on the larger sides, each having a different animal face drawn on it (Figure 2.2c). There are four different objects configurations of growing complexity, each of them used twice, resulting in 8 shapes; (3) A short rectangle box with a one-way mirror glass surface on the top side, which was illuminated from inside when a sensor was touched, using an Arduino circuit (Figure 2.2d,e). Four of the objects were stuck with blue-tac inside the mirror box, arranged at different angles (the position of the objects in the mirror box was determined beforehand and consistent across participants. The orientation of the fixed objects mental rotation by the participant in order to find the right hole); (4) Two paper-printed sensors, made with conductive paint and 3D-printing technology, with a carton hand-shape stuck on top with glue (Figure 2.2d,e); (5) Six Vicon Bonita Motion Capture cameras; (7) Kids' fingerless gloves with attached retroreflective markers.

Figure 2.2

Experimental stimuli and materials, a,b) Object-fitting house, both sides, c) example stimuli set, four complexity levels (low to high: upper left corner to low right corner), d,e) box with light, open and closed, and paper-printed sensors with conductive paint

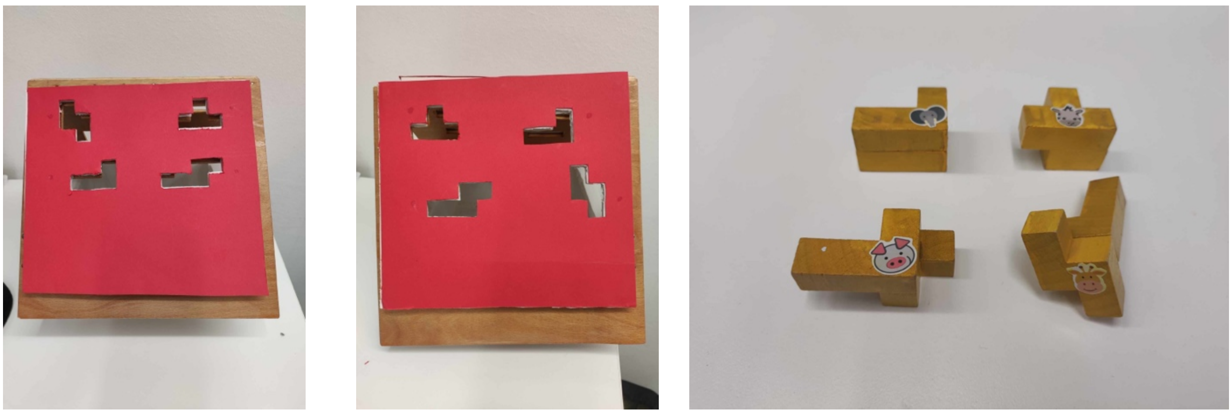

a) b) c)

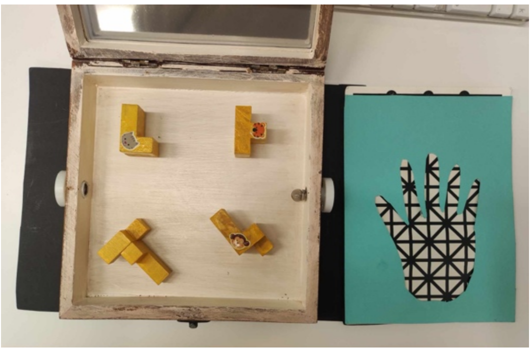

d)

e)

2.2.2.2. Design

Children had to complete all three of the following conditions: In the first condition (*Free Movement Condition*) they were encouraged to manipulate the objects and fit each of them into the correct aperture at the object-fitting toy. In the second condition (*Gesture Condition*) they could see but not interact with the objects, which were inside a box under a glass surface. Their hands were free to point or gesture in any way they wanted while trying to match each shape with the correct hole. In the third condition (*No-Movement Condition*) they were able to see the objects inside the box but they were instructed not to move their hands from certain spots on the table. These spots were hand-shaped surfaces of conductive material, which turned on the light inside the box only when participants kept their hands in touch with the surface (the experimenter kept the light consistently on in the Gesture condition). They then had to state which object fits which hole (Figure 2.3a-c). The conditions were presented with the sequence mentioned above (Free Movement-Gesture-No Movement) for most children (N=8) and randomly for the rest of them (N=16).

2.2.2.3. Procedure

The procedure changed slightly during the pilot phase of the experiment, in an attempt to examine the best sequence and the effect of certain changes on the observed behaviour. For some children (N=6), the following procedure was used: First the child sat at the table and was given the first four shapes to familiarise themselves with. The experimenter then presented the object-fitting house and demonstrated the game and the rules (i.e., that each animal gets in the house through its own gate and no two animals get through the same one, so each one object gets in through a different hole). The experimenter made sure the child understood this constraint by completing a practice trial. The child completed the task alone, while experimenter reminded them of the rules if they tried to put two objects into the same hole and by asking them where each animal goes where next after they completed a trial. When the child completed fitting all objects in one side of the house-fitting toy, the experimenter turned the house around and offered the child the new four objects to fit. When they successfully completed the game, the experimenter presented and positioned the mirror box right between the child's hands, explaining the rules for the next condition (i.e., that they cannot touch the pieces but that they should decide where each animal should go to get inside the house). When they had answered for all 8 objects (i.e., both sides of the toy), the experimenter placed the pad with the sensor under their hands and explained to them that this time they will have to think where each animal must go, but without moving their hands – otherwise the light inside the box would go off, obscuring the objects. When the child answered about all the objects, they earned a reward and proceeded to the working memory game on the tablet.

For the remaining children (N=18), the procedure had the following changes: In the third condition, the light inside the box was kept on by the experimenter and the child had to keep her hands on the table without playing a role in keeping the light on. This choice was made because it became obvious that children were able to keep their hands immobile just by instructing them to do so, and thus the procedure could be simplified. Further, we added different colour codes next to each hole of the object-fitting house, in order for the experimenter to ask about the correct hole without pointing to the apertures – and thus possibly unintentionally encouraging pointing from the child. Finally, a few children (N=5) wore fingerless gloves with Motion Capture reflective markers to measure dynamic kinematics, but such data was only captured from three children.

Figure 2.3

Experimental conditions: a) Free Movement, b) Gesture, c) No-Movement

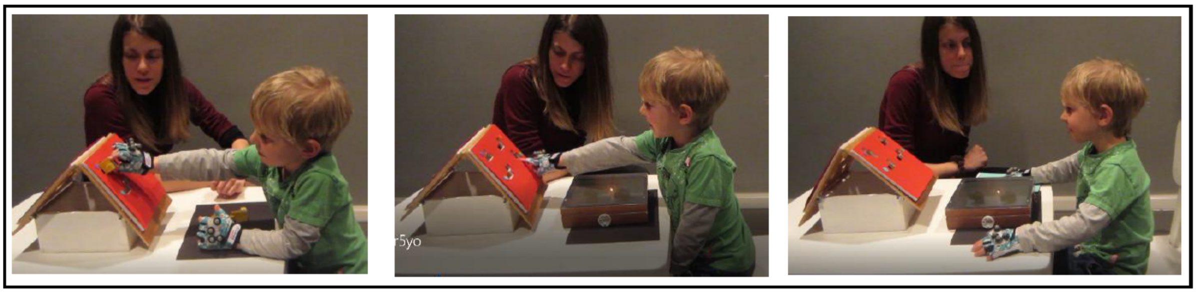

a) b) c)

2.2.3. Visual-spatial working memory task

To measure visual-spatial working memory, we used the “Mr. Ant” task from the Early Years Toolbox (EYT; Howard & Melhuish, 2017). EYT consists of iPad-based direct assessments of early executive function, language, self-regulation, numeracy and social-emotional development. The particular task involves remembering the positions of "stickers" on a cartoon ant and later identifying those positions after a short period of time (Figure 2.4). The difficulty increases gradually as the task progresses, requiring participants to remember the locations of more stickers. Each level of difficulty consists of three trials, ranging from one to eight stickers. The procedure for all trials is as follows: (a) Mr. Ant is presented with a specific number of colored stickers, which corresponds to the current difficulty level, for 5 seconds, (b) a blank screen is shown for 4 seconds, and then (c) an image of Mr. Ant without stickers is displayed, accompanied by an auditory cue prompting participants to recall the sticker locations. Participants respond by tapping on the spatial positions of the stickers on Mr. Ant that they believe were there previously. The task continues until either all three trials at a particular difficulty level are failed or the participant completes Level 8, which involves remembering eight spatial locations. The task was administered using an iPad.

Figure 2.4

Visual WM task trial example: Mr Ant appeared with a sticker on his left cheek. Children had to touch the correct position on a sticker-less Mr.Ant image after a 4s delay between the two presented images

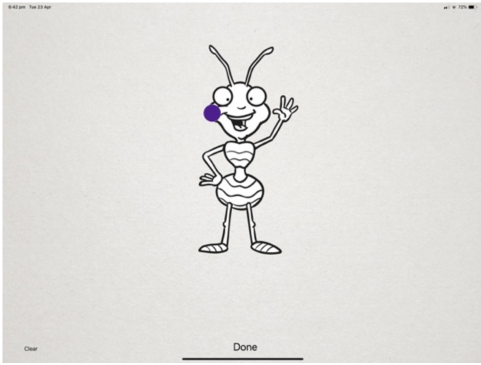

2.2.4. Data collection and coding

Unfortunately, the collection of motion-tracking and EMG data was was made impossible because of the global pandemic and subsequent COVID-19 testing procedures when the labs reopened. Thus we only collected and analyzed behavioural data, which were coded from video recordings. This, as a result, made the detailed analyses of movement kinematics and motor preparation impossible.

We coded the videos using the VLC video player software to identify accurate fitting attempts, hand movement phases during object-fitting, eye fixations in areas of interest and overall time profiles of the problem-solving strategy of the participants. The main variables that were measured were *Fitting Accuracy* in each condition, *Attempts to fit* each object in Free Movement condition, *Fitting time* of each object in Free Movement condition, *Planning Time to Fitting Time* ratio, *Response Time* in Gesture and No-Movement conditions, *Eye fixation time* on box, *Eye fixation time* on house, number of *Fixation Changes* between house and box in Gesture and No-Movement conditions and Pointing in Gesture condition.

In the Free Movement condition, we coded as Fitting Accuracy the number of correct fits the participant made without help by the experimenter, and Attempts to Fit as the number of attempts to insert each object in a hole which was either the wrong hole, or the correct hole with a wrong object orientation., We coded Fitting Time as the time (in miliseconds) that the participants spent to find the correct hole for each object (regardless of whether they were helped by the experimenter or not). Finally, we coded Planning Time as the amount of time spent by participants looking at and manipulating the objects before they started the transport movement towards the holes, and Fitting Time as the amount of time spent to transport and fit the object. We then calculated the ratio, in order to capture the different strategies that were used.

In the Gesture and No-Movement conditions, we coded Fitting Accuracy as the number of correct first responses (or immediate changes of responses from a wrong to a correct one) regarding the hole in which each object would fit in. We coded Response Time as the time in ms which each participant spent to answer about the correct hole for each object. Further, we measured the Eye fixation time in ms that the participants spent in total when focusing on the box and on the Shape-fitting House during their problem-solving phase. However, it worth underscoring that this was a coarse measure as the angle of video capture did not allow for much detail about the specific areas on the box or house that the children fixated. We also measured the number of Fixation Changes between fixating on the box and on the house before each response.

Finally, in the Gesture condition participants only made pointing gestures towards the box and the house as a means of answering the questions asked of them (i.e., there were no representational gestures simulating part of the imagined manipulation or the object). We thus recorded whether the participant pointed to the box and to the house during their problem-solving phase, since pointing might still be a way to offload spatial information and direct the gaze – we were thus expecting it to correlate with WM levels. To code this, we used three levels, namely 0 for no pointing, 1 for pointing to 1-3 objects and 2 for pointing at the box for all objects. Finally, we refer to the objects as Object 1, 2, 3 and 4, reflecting their growing complexity respectively (1 referring to the simplest and 4 to the most complex).

2.3. Results

We first examined the effect of the procedure changes on our data. Wearing the motion capture gloves did not affect accuracies and reaction times in any of the conditions. Moreover, using colour codes and having the participant turn the box light on (Procedure 2) did not affect accuracy, and reaction times either. Since changes in the procedure, motion capture and colour codes did not affect performance, we pooled the data and analysed them as one sample. One participant was excluded from the analyses because part of their video recording was lost.

Accuracy scores and reaction times (both Fitting times and Response times) were tested for normality and differed significantly from normal distributions, so we either conducted non-parametric comparisons (for Accuracy) or used the log-transformed values (for Response Times). Specifically, we conducted Friedman tests as alternative to repeated-measures ANOVA, to check for differences between groups, as well as Mann-Whitney U tests to compare between the two Age groups after a median split (see below). We also calculated the Spearman's rho coefficient for correlations between variables. The main descriptives for each complexity level are reported in Table 2.1.

Table 2.1

Means and SDs for Accuracy and Time spent per Object in each condition

	Accuracy – Free Movement	Fitting Time (ms)	Accuracy - Gesture	Respons e Time – Gesture (ms)	Accuracy – No Moveme nt	Response Time – No Movement (ms)
Complexity 1	0.913 (0.288)	5144.34 (4842.80)	0.708 (0.464)	3529.64 (2041.36)	0.713 (0.657)	4073.88 (3663.84)
Complexity 2	0.783 (0.421)	7642.39 (6515.53)	0.833 (0.381)	3223.59 (1800.57)	0.798 (0.985)	4826.13 (3124.25)
Complexity 3	0.565 (0.507)	17505.83 (16389.17)	0.542 (0.509)	5539.05 (4319.71)	0.493 (0.596)	6499.82 (4898.69)
Complexity 4	0.522 (0.510)	19110.21 (14975.47)	0.543 (0.592)	4340.45 (2559.40)	0.477 (0.558)	2977.11 (1080.56)

The participants' Age did not significantly affect their accuracies in any condition (Free Movement: $\chi^2(2) = 2.769$, $p = .250$, Gesture: $\chi^2(2) = .447$, $p = .800$, No Movement: $\chi^2(2) = 1.341$, $p = .512$). In an attempt to reveal developmental differences within the age range we examined, we split the children (based on median age) into Younger and Older participants. The effect of Age Group was still not significant (Free Movement: U = 91.5, p = .221. Gesture: U = 56, p = .326, No Movement: U = 67.5, p = .786).

Moreover, there were no statistically significant differences in accuracy between the different movement conditions ($\chi^2(2) = 2.032$, $p = .362$) either.

The four objects differed significantly in number of attempts to fit in the holes in Free Movement condition ($\chi^2(3)=24.387$, $p <.001$) (Figure 2.5). Post-hoc analyses with Wilcoxon signed-rank tests showed significant differences between the attempts to fit Object 1 compared to Object 3 (Z = -3.264, p = 0.001) and compared to Object 4 (Z = -3.253, p = 0.001). Similarly, attempts to fit Object 2 differed significantly from attempts to fit Object 3 (Z = -2.857, p = 0.004) and Object 4 (Z = -3.041, p = 0.002). On the other hand, attempts to fit the two simplest objects did not differ significantly (Z = -1.403, p = 0.161) and neither did the attempts for the two more complex ones (Z = -0.076, p = .939).

Figure 2.5

Number of fitting Attempts per Object in Free Movement Condition

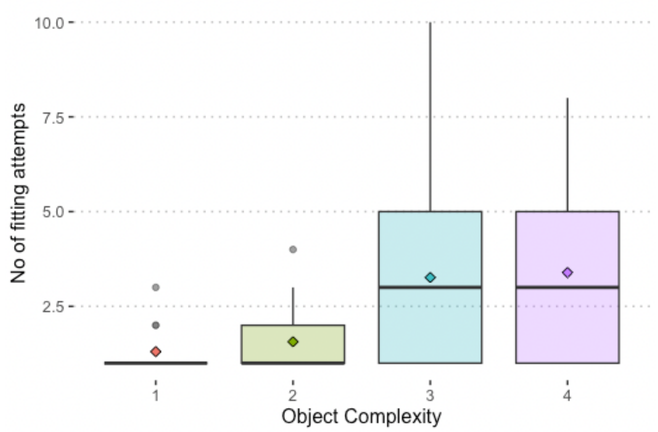

The Fitting Time between the objects in Free Movement condition also differed significantly (F(3)=9.708, p < .001) (Figure 2.6a). Specifically, Object 1 took significantly less time to fit than Object 3 (p = .007), and Object 4 (p < .001). Also, Object 3 took significantly less time to fit than Object 4 (p = .006). There was a significant difference in Response Times for each object in Gesture condition (F(3) = 2.831, p = .045). Post-hoc tests however did not show significant differences between

Response Times for each of them (Figure 2.6b). In the No Movement condition, Response time again differed significantly ($F(3) = 7.059$, $p < .001$; Figure 2.6c) between objects. Specifically, Object 3 took significantly more time to fit than Object 1 ($p = .024$) and Object 4 ($p = .002$). The fact that Object 4 did not have the greatest difficulty to fit in No Movement condition, was probably because it was chosen last by the children to fit, after answering for the previous objects – it was thus a restrained choice.

Figure 2.6

Fitting/Response time per Object in a) Free Movement, b) Gesture, and c) No Movement condition

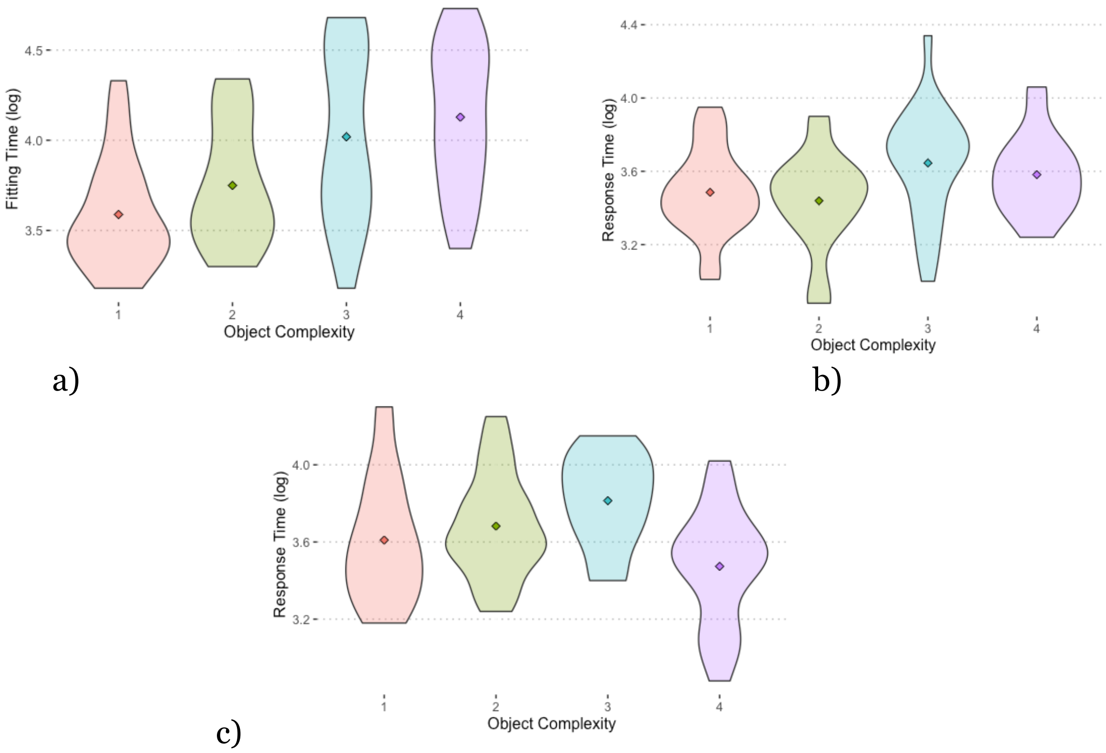

Furthermore, the Planning to Fitting time ratio did not correlate significantly with Age ($r=0.076$, $p = .366$), but it positively correlated with Accuracy in Free Movement ($r=0.345$, $p = .05$), suggesting that spending more time to plan the transport and rotation produced more accurate fittings (Figure 4a).

Working memory levels positively correlated significantly with Age ($r=0.607$, $p = .001$) and showed a trend but did not significantly correlate with Accuracy in Free Movement ($r=0.304$, $p = .085$), neither with Accuracy in Gesture ($r=0.064$, p=.38), nor with Accuracy in No Movement condition (r=0.162, p = .22). Finally, the amount of Pointing in Gesture condition was not associated with Age (r=-.266, p=.11) nor with Working Memory (r=.132, p=0.56).

2.4. Discussion

Our findings confirm our first hypothesis that preschoolers would find it more difficult to fit the more complex objects in the Free Movement condition, needing more time and more attempts to achieve this. However, the difference between the two simpler and the two more complex objects was not systematic – children needed more time to fit Object 4 than Object 3, but they made an equal number of attempts. Similarly, they spent a similar amount of time to fit Object 1 and 2. The fact that the object complexity did not produce a linear increase in difficulty can also be observed in the Gesture condition, where response time differences between the simpler objects and Object 3 was marginal – and non-significant in the pairwise comparisons. A consequence of our design in the Gesture and No Movement conditions made it impossible to actually measure the response time for Object 4 meaningfully (it was always last and so the possible solutions was much more highly constrained than for the other objects) so it is difficult to draw conclusions about it. That said, this lack of consistency in the effects of object complexity observed is also probably explained by the exact geometrical features we manipulated in their design: Object 1 and 2 differed in their symmetry but they could be fitted in the same number of ways, while Object 3 and 4 had their complexity manipulated in a non-comparable way: either across symmetry, or across the axis of elongation which should be used to fit the aperture. To better control and measure such effects, these characteristics should be manipulated separately, as it has been shown in previous studies (Fragazcy et al, 2015; Ornkloo & vanHofsten, 2007; Street, James, Jones, & Smith, 2011).

Moreover, we did not find significant age differences in terms of accuracy, suggesting that, although the different skills relevant to object fitting are still developing during our age range, and different children are more or less proficient, all children eventually manage to fit the shapes successfully given enough time. Since we were mostly interested in the fitting process, we did not use a stricter measure of accuracy (e.g., successful pre-adjustments, as in Ornkloo & vanHofsten, 2007). When looking into the fitting process, a difference between more and less efficient fitters was evident and had a significant correlation with the time that the children spent planning (by looking and manipulating the object) before they decide to transport it towards a hole. This different strategy could be explained by a lower ability to mentally represent and remember many spatial relations between object parts (e.g., asymmetrical shapes need larger amounts of information to be processed and stored) and, thus, a lower ability to mentally rotate. As a result, when this offline capacity is lower, the environment has to be exploited more in order to support the mental processes – in this case, more actual comparison of the object and the aperture until they are aligned.

Another interesting – and possibly related - observation in the Gesture condition, was that children used pointing gestures to guide their attention to relevant items and to highlight their important characteristics as they attempted to choose the correct hole for each object, although they did not use representational gestures. In this case, hands were possibly used as a stable mnemonic anchor in the environment that allowed them to move their eyes less between the possible solutions.

Such connection between pointing and eye movements has been reported before (Cappuccio, Chu & Kita, 2013). However, we did not find a relationship between the pointing behavior and WM. This might be explained by the different use of pointing by children, both as a communicative gesture to answer questions and as a mnemonic support – these two types of pointing were not differentiated in our coding scheme and could thus have obfuscated the actual effect. To actually tap into attentional processes and how hand movements are possibly used to direct the eyes (an hypothesis opposite or complimentary to visually-guided action which we previously discussed) would need a detailed real-time capture of eye and hand movements during real and imagined object manipulation. Unfortunately, while planned, this was not possible because of the Covid pandemic and the associated testing restrictions.

Finally, since accuracy was not compromised when hand movements were prohibited (No Movement condition), we conclude that, while useful in directing attention, gestures were not necessary for children to solve this particular task. However, we should emphasise that most studies that report a benefit in performance from co-thought gestures in mental problem-solving tasks explicitly encourage people to use gestures to help themselves (e.g., Chu & Kita, 2011), whereas we expected this behaviour to be expressed in a spontaneous way.

In general, our experimental design would benefit from more nuanced manipulation of object complexity across different attributes, along with measures which will reveal eye and hand movements during fitting and mentally rotating objects – as well as possible standardised measures of mental rotation skills. This would allow as to record aspects of the fitting process in more detail, as well as the interplay between action and vision, specifically in the context of using the body while sampling and processing information in the environment.

In the next chapter we move to investigating the value of information for different age groups, and we explore how school-aged children, adolescents and adults balance their preferences between gaining external rewards, resolving their uncertainty and exploring novel stimuli. We will return to an examination of object complexity and its effect on preschoolers' exploration in Chapter 4.

Chapter 3

Real-time planning of explorative and exploitative actions across development: Four mouse-tracking studies

3.1. Introduction

From a very early age humans act intentionally on their environments in order to reach a desirable state (e.g, obtain a goal). This desired goal can be an internal or external state and can be reached in a shorter or longer time. Everyday life is a constant balancing act or struggle between these two types of goals. Goal-directed actions, as a psychological process, have been predominantly investigated in instrumental contexts, where an agent aims to change the state of the world and success is externally rewarded. However, humans constantly also plan epistemic actions, where they aim in changing their internal knowledge state (i.e., they actively interrogate their environment to gain information and improve their understanding of the world). These actions are also rewarded but by experiencing the intrinsic value of the learning progress, rather than obtaining a clear external reward such as food.

As both types of rewards (internal and external) motivate actions, humans often experience conflict between choosing a known, rewarding option and the option to try something new (e.g., a new bar, a new holiday destination, a new toy). This situation is known as the *exploration-exploitation* dilemma (Sutton & Barto, 2018; Cohen, et al., 2007; Melhorn et al., 2015), This dilemma highlights the following problem: sticking to the familiar option will yield immediate and predictable rewards, but might never lead to more rewarding options, whereas trying the new option might or might not lead to external reward, but it will lead to learning and will improve future decisions. The optimal solution to this dilemma is dependent on the specific demands of the decision context. For instance, increased exploration is optimal when the agent has less knowledge about the environment or when the environment changes frequently and has large variability, whereas exploitation is better suited for known environments that remain relatively stable. Thus, the best choice might depend on context and also perhaps on the level of development (e.g., children might benefit from broader exploration in general, as their knowledge of the world is smaller).

In fact, children have indeed been shown to explore more in uncertain situations (i.e., where information is ambiguous and contradicting or when prior knowledge assumptions are violated; Bonawitz, van Schijndel, Friel, & Schulz, 2012; Cook, Goodman, & Schulz, 2011; Schulz & Bonawitz, 2007; Taffoni et al., 2014; vanSchijndel Visser, van Bers, & Raijmakers, 2015). More specifically, children seem to keep exploring while there is still learning progress and uncertainty in their environment (Liquin, Callaway, & Lombrozo, 2021: Ruggeri, Pelz, Gopnik, & Schulz, 2021). Similar findings have been reported with infants (Addyman & Mareschal, 2013; Chen, Westermann, & Twomey, 2022; Kidd & Haynes, 2015; Kidd, et al., 2014; Poli et al., 2020; Sim & Xu, 2017).

However, although common experience and theory have suggested that infants and young children are more explorative than adolescents and adults, or at least that their exploration is broader (Gopnik, 2017; Sumner et al., 2019), research has only recently started to shed light on the exact differences in their exploratory strategies. Specifically, it has been suggested that children might be different in the relevant amount of directed and random exploration they engage in when searching a new environment (Gopnik, 2017; Wilson et al., 2021). Since directed exploration is considered a more sophisticated ability, it has been suggested that it could be related to the development of cognitive control (Blanco, et al., 2015; Badre, Doll, Long, & Frank, 2012; Otto, Knox, Markman, & Love, 2014). However, the dissociation between directed and random exploration is based on the presence vs. absence of sufficient knowledge to guide behaviour, while different informational attributes might also be relevant for exploratory behaviour, and might possibly be weighted differently in different ages. Specifically, it could be the case that novelty might influence exploration in a dissociable way compared to uncertainty, reflecting a broader possibility for learning. In this case, it would stand somewhere between directed exploration which focuses specifically on missing parts of information, and random exploration, which reflects a complete lack of knowledge, i.e., if one doesn't know the specific source for uncertainty resolution, choosing a new option will always facilitate learning. The separate influence of novelty on learning is documented since infancy (e.g., Poli et al., 2020), but the relative influence of uncertainty and novelty across development is still not extensively studied.

Several recent studies have directly tried to assess the differences in exploratory strategies in children, adolescents and adults. For example, in a study by Schulz et al. (2019), children in middle and late childhood and adults completed a spatially correlated multiarmed-bandit task. Children engaged in more directed exploration and generalised less than adults, but they did not show any differences in the amount of random exploration. Similarly, Meder et al. (2021) used a similar task to study children's exploration strategies at 4 and 9 years of age, finding that random exploration decreased as children grew older. They also found evidence of directed exploration in the youngest group. Somerville et al., (2017) studied 12 to 28 year-olds, showing that adolescents were strongly driven by immediate rewards. Jepma, Schaaf, Visser and Huizenga (2020) compared adolescents' and adults' exploratory strategies and learning rates and showed that, overall, adolescents explore more and assume greater environmental volatility compared to adults. Recently, in an attempt to dissociate between the influence of novelty and uncertainty, Nussenbaum et al. (2022) studied participants aged from 8 to 27 years in an exploration task that separately manipulated stimulus novelty and reward uncertainty. They found that children explored the uncertain option more than older participants, but risk aversion was probably a strong factor which influenced adolescents' and adults' behaviour. Overall these findings point towards certain directions but show that the relationship between maturation and the exploration-exploitation balance is still unclear.

Taking these open questions into account, we designed and conducted a series of studies that aim to address some aspects of the exploration-exploitation dilemma across development. The studies explore the ongoing conflict between possible behavioural options during a choice decision process, and its relation to cognitive control. We designed a decision-making task in which participants had to choose between three possible options across a number of trials: (i) an externally rewarding option (hereafter called ER), (ii) an option revealing missing information (hereafter called IR, informational reward) and (iii) an option offering random novel stimulation (hereafter called NS). The task differed in some aspects from a classic explore-exploit task (e.g., n-armed bandits; Sutton & Barto, 2015; Witten, 1977) in several ways. First, we were interested in dissociating informational value from the acquisition of the reward, such that learning could be motivating per se, regardless of one's desire to obtain a reward. Second, we were interested in both short-term and long-term rewards (both external and intrinsic) because real-life situations often involve persistence in the face of other valuable distractors to achieve a reward (e.g., as in delay discounting tasks; da Matta, Goncalves, & Bizarro, 2012). Previous studies have also shown that longer temporal horizons make exploration more valuable (Somerville et al., 2017; Wilson et al., 2014). Finally, our task had pre-established and stable external rewards, which the participant explicitly learned from the training phase, in order to avoid exploration and curiosity relevant to the reward (i.e., they knew exactly what they should choose to obtain the reward, and what this reward would be).

We were also interested in investigating the role the development of cognitive control on the selection of a specific exploratory strategy. To this end, we drew from a different research line suggesting that subtle differences in hand movements can reveal differences in participants' cognitive control during real-world online decision making. For example, control processes such as inhibition and attention can be revealed by tracking hand movements in a 3D space (in adults: Erb, Moher, Sobel, & Song, 2016, and in children: Erb, Moher, Song, & Sobel, 2017, 2018), as well as mouse movements in computerised tasks (Benedetti, Gronchi, Gavazzi, Bravi, Grasso, Giovannelli, & Viggiano, 2021; Dieciuc, Roque, & Boot, 2019). Mouse-tracking is widely used to assess ongoing decision processes, and often also to reveal preferences in value-based decision making (Koop & Johnson, 2013; O'Hora et al., 2016). Because our task involved comparisons between subjective values of different options, this seemed an appropriate way to reveal differences in individuals' subjective evaluations that might be overlooked by simply focusing on their final choices.

Specifically, based on the literature reviewed above, we anticipated differences between age groups in their choice preferences in the decision-making task, such that the child group would choose the NS option more than the other two groups, who would not differ in terms of preferences for the NS option. We also expected adolescents to choose the ER option more than the other options. Furthermore, we expected to find differences in their mouse movements, such that the children's movements would reflect greater conflicts between the ER and IR options, as well as the IR and NS option than either adolescents or adults. More specific hypotheses in each experiment are discussed separately below.

3.2. Experiment 1

In our first experiment, we compared children from 5 to 9 years of age, adolescents aged 13 to 16 years and adults up to 35 years. The age range was chosen to reflect the age ranges of similar experiments (e.g., Wu et al., 2019). Pilot studies also revealed that 5 years of age was a lower bound of children who could (i) use a mouse efficiently, (ii) complete the task, and (iii) not produce too much extreme movement noise.

The decision-making task involved choosing between three options, but in a two-alternative forced choice format per trial (2AFC; i.e., participants had to choose one of the two presented options for the experiment to proceed). This allowed for direct comparisons between different dilemmas. In this initial task, each experimental block ended when the participants obtained the ER, or when a maximum number of 27 trials elapsed, allowing participants to win the ER without completing the total number of the block trials (e.g., by consistently making hard ER-driven choices). This strategy provided a short temporal horizon and could thus lead to less exploration if participants were very reward-oriented.

We expected all groups to obtain the ER in the majority of the blocks, but to show differences in the frequencies of their choices of IR and NS. We also expected all groups to choose IR more than NS, but children to choose NS more than adolescents and adults, sacrificing some of their IR choices to explore NS. In terms of movement parameters, we anticipated differences in temporal parameters and geometrical-complexity parameters in different dilemmas in all age groups. Specifically, we expected more conflict (i.e., longer times, later commitment, more deviation, more entropy) in ER-IR and IR-NS dilemmas for children, whereas we anticipated more conflict in ER-IR for adults. Adolescents might also be conflicted in ER-IR, but individual differences might be large.

3.2.1. Methods

3.2.1.1. Participants

Our participants consisted of 23 children (5 females, mean age: 7.09 years), 20 adolescents (12 females, mean age: 14.47 years) and 20 adults (11 females, mean age: 32.5 years). All participants were neurotypical and had normal or corrected-to-normal vision. The participants were volunteers recruited mainly from word-of-mouth and though the Birkbeck Babylab database. They received a £5 Amazon voucher for their participation.

3.2.1.2. Design

Participants in all age groups had to complete three experimental blocks of a maximum of 27 trials each. In each trial, two of the total three options (selected randomly) appeared on the screen. The participants had to choose an option in a 2AFC format for one of the following dilemmas: External Reward vs. Informational Reward (ER-IR), External Reward vs. Novel Stimulation (ER-NS), or Informational Reward vs. Novel Stimulation (IR-NS). When the External Reward (ER) option has been chosen for nine times in total, the participant won the reward and the block ended. Otherwise, the trials continued until ER was acquired or a total of 27 trials was completed. This means that the maximum number of choices of the other curiosity options could be for a maximum of 18 for each type of curiosity (total number of trials (27)/3 options = 9, presented by 2 every time, i.e., 18 times for each option), whereas the ER option could only ever be chosen for a maximum of 9 times within a block.

3.2.1.3. Stimuli and procedure

Due to the COVID-19 pandemic measures, the task was delivered online, through the participant's own home computer using a computer mouse. The use of a laptop's touchpad was not encouraged but it was allowed. Participants could also use a tablet. The task was presented as a game in which participants had to make decisions between three dogs and gain a reward at the end of every block, based on their choices. The choice was operationalised by clicking on a dog bone and dragging it towards one of two possible dogs. (Figure 3.1). The reward strategy was explicitly stated from the beginning: They were told that choosing one of the dogs (the red one) led to the accomplishment of the main goal: the construction of a tower – but they were not explicitly encouraged to do so. However, the other two options were also designed to be attractive, although not explicitly rewarded.

Figure 3.1

The different dilemmas (ER-IR, ER-NS, IR-NS).

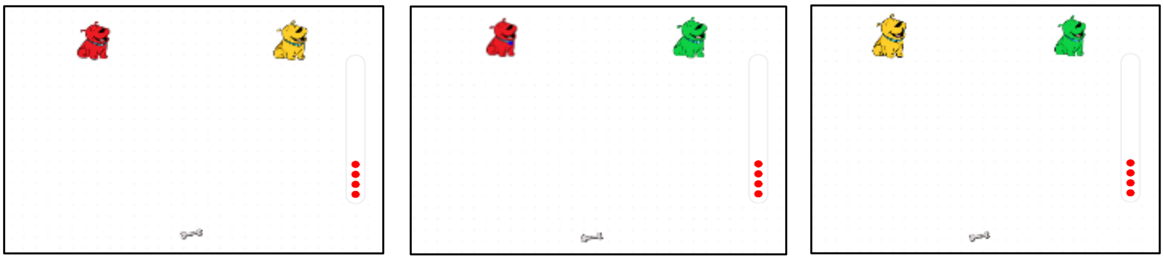

More specifically, the game unfolded as follows:

(i) Training trials: One of the three dogs (red, green or yellow) appeared in random positions at the top of the screen. Every dog appeared for 9 trials. The participants had to drag the bone from the bottom middle starting location and drop it on the dog. When they achieve this, the screen changed to one of three possible situations (see Figure 3.2). When the red dog was chosen (an *ER* choice), a screen with a tower appeared. Every time this option was chosen, another block was added onto the tower until it reached a cloud. When the yellow dog was chosen (an *IR* choice), a screen with a puzzle appeared. Every time this option was chosen, another piece of the puzzle was revealed. When the green dog was chosen (an *NS* choice), a screen with a cartoon character appeared. Every time this option was chosen, a different character would appear.

Figure 3.2

The effects associated with each choice: Red Dog: Build the tower with increasing blocks (main goal – Externally Rewarded), Yellow Dog: Reveal the puzzle with increasing pieces filled (Informational Reward), Green Dog: See a new image (Novel Stimulation)

(ii) Experimental trials: two of the three dogs (pseudo-randomly selected) appeared in one of two fixed positions each, at the top of the screen. The participants had to drag the bone to one of the dogs, thereby making a choice. A bar with the remaining lives to the end of the block could be seen on the right side of the screen, decreasing one at a time with every choice (Figure 1). Depending on the choice, the relevant second screen appeared, as in the training set. Whenever the participant managed to complete the tower (by choosing the red dog 9 times), the block ended and a WIN screen appeared, followed by a screen which showed the rewards gained. Otherwise, trials continued until a maximum number of 27 was completed. In this case, a TRY AGAIN screen appeared. Children completed a slightly different version of the task than the one for the adolescents and adults. Specifically, the puzzles in the child version showed scenes from the popular cartoon movie "Monsters Inc." (Figure 3.3a) and the character images were animal cartoons (cats, penguins and jungle animals; Figure 3.3b). In contrast, in the adolescent and adult version the puzzles showed surrealistic collages (Figure 3.3c) and the character images were animal photographs (cats, dogs and horses; Figure 3.3d). Furthermore, the children's rewards were images of donuts and a maximum of three could be gained, one for each block of trials. The adult and adolescent rewards were images of shop vouchers, and a maximum of three could be gained as well. In both cases the promise of winning the actual rewards (sweets and vouchers) was made.

Figure 3.3

Examples of IR and NS stimuli for children's and adolescents-adults' versions.a) Children IR, b) children NS, c) adolescent-adult IR, d) adolescent-adult NS

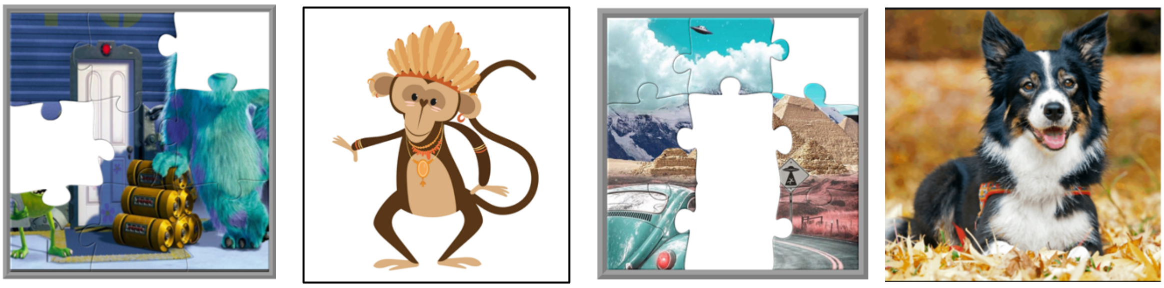

a) b) c) d)

The participants were contacted via video-calls and, after some familiarisation with the experimenter, they were asked for consent. They were then presented with the training trials. After completing this, the instructions of the main game were explained to them. They were given the time to ask any relevant questions that they may have about playing the game. After each set of experimental trials (a block), the participants were reminded of the instructions. More specifically, they were reminded that the puzzles and characters would be novel for the new block.

3.2.2. Results

No participants were excluded from this experiment, although there were three individuals (2 children and 1 adult) who did not complete the final block of trials due to technical issues with their computers. We first describe the choice data and then report the motion parameter analyses.

3.2.2.1. Choices analyses

Table 3.1 shows the main descriptives for the different choices made by each age group. We first compared how many times each group chose each option with a two-way ANOVA and found a significant interaction between Age Group and Option ($F(3.612,120) = 2.706$, $p = .039$, Greenhouse-Geiger corrected). Post-hoc Bonferroni-corrected comparisons showed that in total, participants chose the ER option more than the IR option ($p<.001$) and the IR option more than the NS option ($p<.001$) (Figure 3.4a). Simple effects revealed that adolescents chose the ER option significantly more than adults ($p = .041$), but not children ($p = .702$), while children and adults also did not differ in their ER choices ($p = .513$). Furthermore, groups did not differ in their IR choices, but adults chose the NS option significantly more than adolescents ($p = .026$). There was no significant difference between the children and adolescents ($p = .998$) and children and adults NS choices ($p = .208$). Within each group, children chose the ER option significantly more often than the IR option ($p < .001$) and the NS option ($p < .001$). They didn't differ in their IR and NS choices ($p = .059$). Adolescents chose ER significantly more often than the IR option ($p < .001$) and the NS option ($p < .001$). They also chose the IR option more often than the NS option ($p = .007$). Adults only had a marginally significant difference in their ER and NS choices ($p = .051$). Their ER and IR choices did not differ ($p = .374$), neither did their IR and NS choices ($p = .210$).

Table 3.1

Means and SDs of choices per Option for each age group

	Children	*Adolescents*	*Adults*
External Reward	8.287 (1.050)	8.842 (0.688)	7.939 (1.605)
Informational Reward	5.515 (2.260)	5.280 (2.480)	6.787 (2.793)
Novel Stimulation	3.893 (3.197)	3.017 (1.995)	5.257 (3.733)

Figure 3.4

a) Differences between Option choices across groups, b) The younger adolescents group had no significant difference in the number of IR and NS choices.

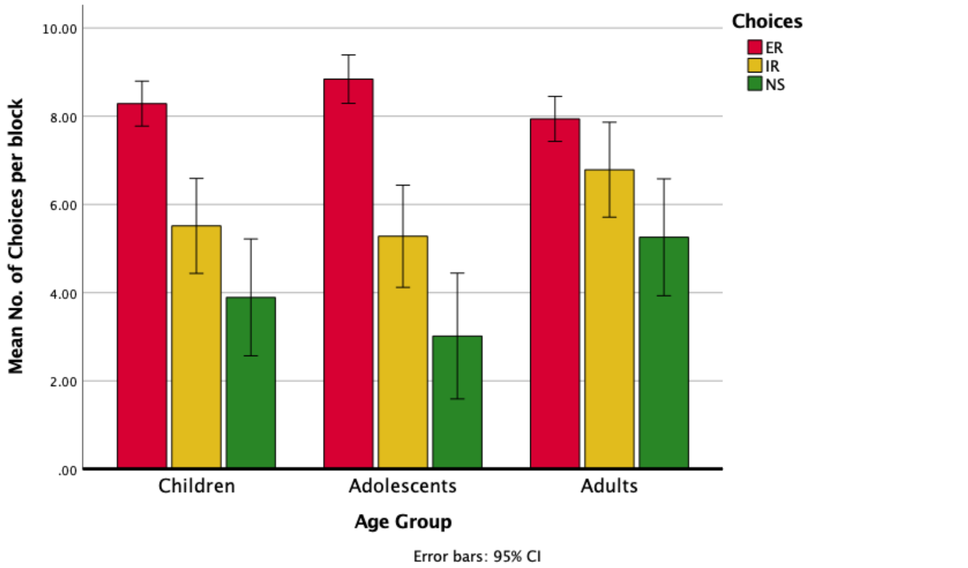

Since the children and adolescent group had a large age range, within which developmental differences in decision-making might occur, we split these two groups into two each and examined their choices (Figure 3.4b). The two child age groups (5-7.5 year-olds, 7.5-9 year-olds) did not differ in terms of choices ($F(1,21) = .054$, $p = .818$) and had the same pattern of significant differences between all options.

However, there was a significant interaction between membership in the adolescent age groups (13-14 year-olds, 15-16 year-olds) and choice ($F(2,34) = 5.440$, $p<.01$). Specifically, the younger adolescent group had no difference in their choices between the IR and NS options, similar to the children group.

3.2.2.2. Mouse-tracking analyses

Pre-processing of movement data

To analyse the data we collected through mouse-tracking, we broadly followed the process discussed by Wurff et al. (2021). To preprocess the data and calculate the variables of interest we used the package *mousetrap (http://pascalkieslich.github.io/mousetrap/),* which was built by Wurff and colleagues in the R programming language (R core team, 2020). Further statistical analyses were also performed using R. The *Mousetrap* package allows for the calculation of the most commonly used parameters in decision-making experiments; namely, trajectory indices and temporal measures, while offering some more sophisticated analysis options regarding types of trajectories that are commonly observed in such experiments. In our experiment, although we calculated a range of several measures, we will be reporting the analyses of five trajectory indices: three curvature measures, *Maximum Absolute Deviation*(MAD), *Maximum deviation above the ideal line* (MDabove; i.e., taking into account only the deviation towards the alternative option) and *Area under the Curve*(AUC); one complexity measure (*x-flips*); and one temporal measure, *Response Time (RT)*). Furthermore, we fitted the trajectory data to predefined (by the mousetrap package) *trajectory types* which are believed to reflect different cognitive processes such as discrete vs. continuous decision-making process (Wulff, Haslbeck, Kieslich, Henninger, & Schulte-Mecklenbeck, 2019). We observed instances of all the proposed trajectories: Straight, Curved, Continuous Change of Mind (cCoM), Discrete Change of Mind (dCoM) and Double Discrete Change of Mind (dCoM2). We then compared the frequency of these types in each dilemma.

To calculate these parameters, we spatially and temporally normalized our raw data (for trajectory and temporal measures respectively), and filtered out any anomalous trajectory that deviated by more than 2 SDs from the different prototypes already included in the mousetrap package. These trajectories are basically the ones with "erratic output producing non-interpretable looping cycling leftward and rightward" (Freeman et al., 2008).

Analyses

Table 3.2 shows the main descriptive values for all movement parameters and Table 3.3 depicts the frequency of trajectory types per dilemma and Age Group. We used R (R Core Team, 2020) and *lme4* (Bates, Maechler & Bolker, 2012) to perform a mixed effects analysis of the relationship between each parameter and the three choice environments (*Dilemmas*: ER-IR, ER-NS, IR-NS). We used the log-transformed values for RTs, as these violated the normality assumption (a common case with response times data, Lo & Andrews, 2015). Moreover, we built generalised mixed effects models for MDabove, using a Gamma distribution with a log link function, and similarly for xflips, using a Poisson distribution with a log link function. To compare the frequencies of trajectory types per Dilemma and Age Group, we performed ordinal regressions, using a cumulative link mixed model. For all the measures, as fixed effects, we entered Dilemma and Age Group and their interaction into the models. To incorporate the dependency among observations of the same subject and dilemma, as random effects, we had intercepts for subjects and by-subject random slopes for the effect of Dilemma. For example, the model for RTs including all age groups was the following:

RT.model = RT ~ Dilemma*AgeGroup + Dilemma + AgeGroup + (1+Dilemma|subject)

P-values were obtained by restricted maximum likelihood ratio tests of the full model with the effect in question against the model without the effect in question (chi-square tests for nested models).

Table 3.2

Means and SDs for movement parameters per Dilemma and Age Group

		Children		Adolescents		Adults	
Features[a]		Mean	SD	Mean	SD	Mean	SD
Maximum Absolute Deviation (MAD)	ER-IR	40.24	139.84	21.45	121.84	19.06	118.34
	ER-NS	21.52	122.30	17.33	101.45	19.32	112.02
	IR-NS	54.62	165.12	34.91	127.19	44.43	156.06
Maximum Deviation above ideal line (MDabove)	ER-IR	66.92	121.24	47.25	106.55	48.65	101.06
	ER-NS	53.61	102.90	42.63	84.24	48.06	93.83
	IR-NS	86.79	137.80	57.55	112.33	69.69	140.39
Area Under the Curve (AUC)	ER-IR	10378.46	44327.26	6125.90	41060.39	2946.68	32365.41
	ER-NS	4902.83	35770.59	3882.43	32221.04	2944.88	29655.73
	IR-NS	14984.84	45897.18	8939.75	39801.28	8593.42	41403.56
x-flips	ER-IR	1.09	1.46	0.51	0.79	0.70	0.99
	ER-NS	0.96	1.13	0.51	0.84	0.77	0.91
	IR-NS	1.21	1.58	0.57	0.85	0.76	1.15
Response Times	ER-IR	1340.56	1041.81	927.63	534.11	904.53	491.15
	ER-NS	1347.69	1018.71	974.56	631.35	900.37	515.89
	IR-NS	1424.55	1030.61	966.09	571.01	950.44	601.66

[a]All time related values are presented in milliseconds (ms), all position related values are presented in pixels (px), area (AUC) is displayed in px².

The unfolding of aggregated trajectories across time steps for each Dilemma can be seen at Figure 3.5. There was no significant DilemmaXAge Group interaction effect on MAD ($\chi^2(4) = 1.991$, $p = .738$). However, the type of Dilemma significantly affected MAD ($\chi^2(2) = 28.192$, $p < .001$). Specifically, there was greater deviation in the IR-NS dilemma than the ER-NS ($p = .003$) and the ER-IR ($p = .022$) (Figure 3.6a). There was no significant difference due to Age Group ($\chi^2(2) = 1.913$, $p = .384$). No significant interaction between Dilemma and Age Group was observed for

MDabove ($\chi^2(4) = 2.084$, $p = .720$), but type of Dilemma significantly affected MDabove ($\chi^2(2) = 32.986$, $p <.001$) (Figure 3.6b). There was greater deviation in the IR-NS dilemma than the ER-NS ($p <.001$) and the ER-ER ($p <.001$). Age Group did not have a significant effect ($\chi^2(2) = 2.220$, $p = .330$). Furthermore, there was no significant DilemmaXAge Group interaction effect on AUC ($\chi^2(4) = 2.168$, $p = .705$). Type of Dilemma had a significant effect on AUC ($\chi^2(2) = 11.773$, $p = .003$), with the IR-NS dilemma having larger AUC than the ER-NS ($p = .003$) (Figure 3.6c). No significant effect was observed due to Age Group ($\chi^2(2) = 3.566$, $p = .168$). There was no significant DilemmaXAge Group interaction effect on xflips ($\chi^2(4) = 6.888$, $p = .141$). No main effect of Dilemma was observed ($\chi^2(2) = 2.385$, $p = .304$), but there was a main effect of Age Group ($\chi^2(2) = 16.653$, $p < .001$). Specifically, children did more xflips than adolescents ($p<.001$) and adults ($p = .021$). Last, no significant DilemmaXAge Group interaction effect was observed on RTs ($\chi^2(4) = 1.290$, $p = .863$). There was a main effect of Dilemma on RTs ($\chi^2(2) = 6.269$, $p = .043$) (Figure 3.6d). Specifically, participants were slower at the IR-NS dilemma than the ER-IR dilemma. There was also a significant effect of Age Group on RTs ($\chi^2(2) = 11.850$, $p = .003$). Specifically, children were slower than adolescents ($p = .030$) and adults ($p = .007$).

Figure 3.5

Change of mouse position on x-axis during the progress of each trial (time-normalised trajectories). Data have been aggregated for each dilemma. A slower unfolding of the decision at the IR-NS dilemma can be observed in the full sample data (a). Different unfolding of decisions over time can be observed in children (b), adolescents (c) and adults (d).

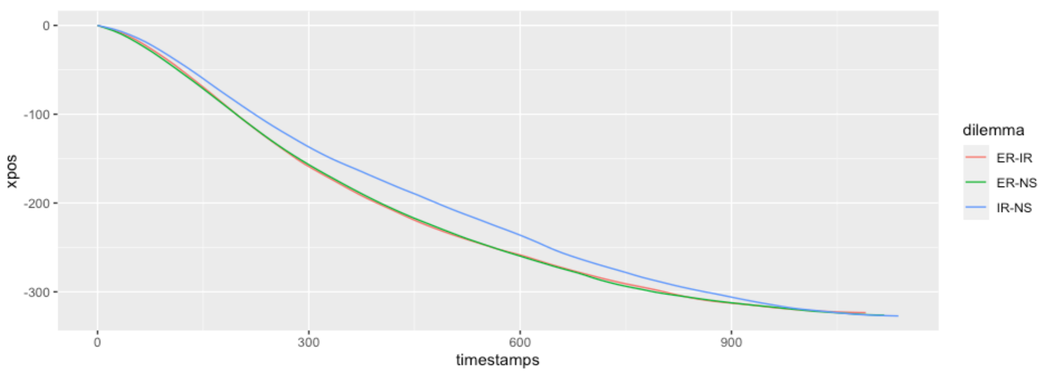

a)

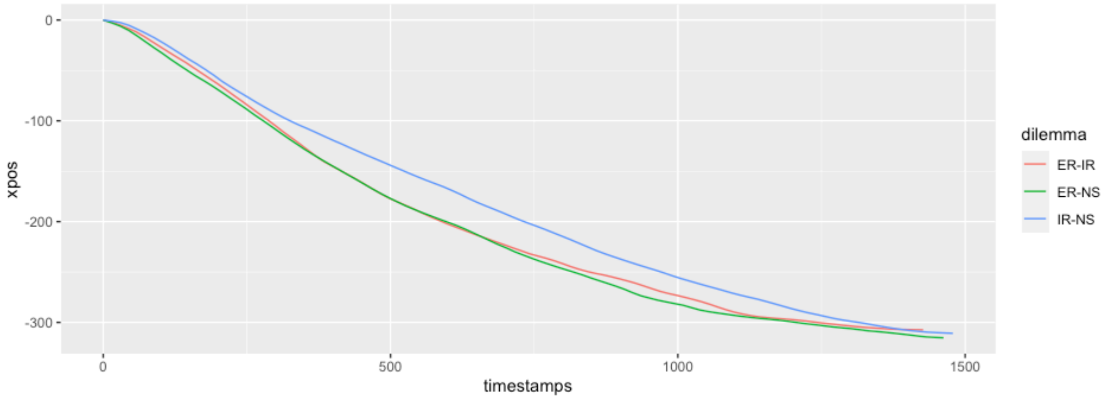

0
-100
-200
-300
xpos
0
500
1000
1500
timestamps
dilemma
ER-IR
ER-NS
IR-NS

b)

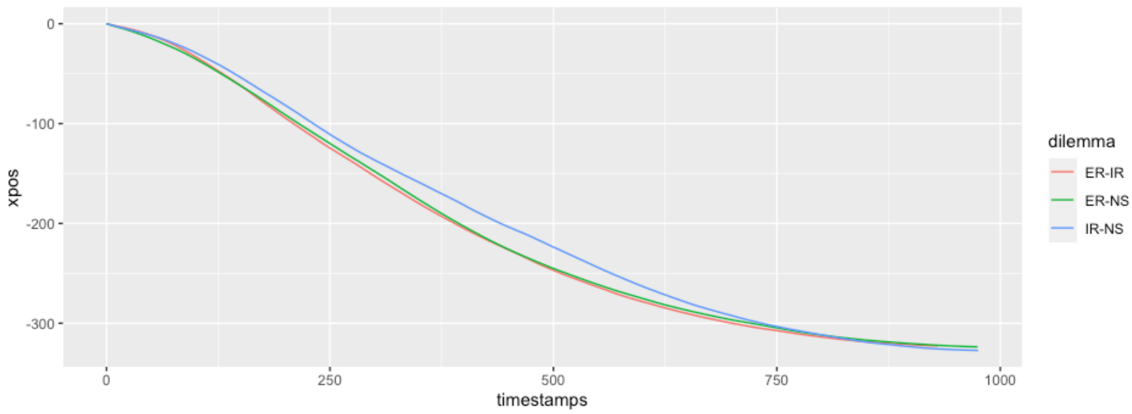

0
-100
-200
-300
xpos
0
250
500
750
1000
timestamps
dilemma
ER-IR
ER-NS
IR-NS

c)

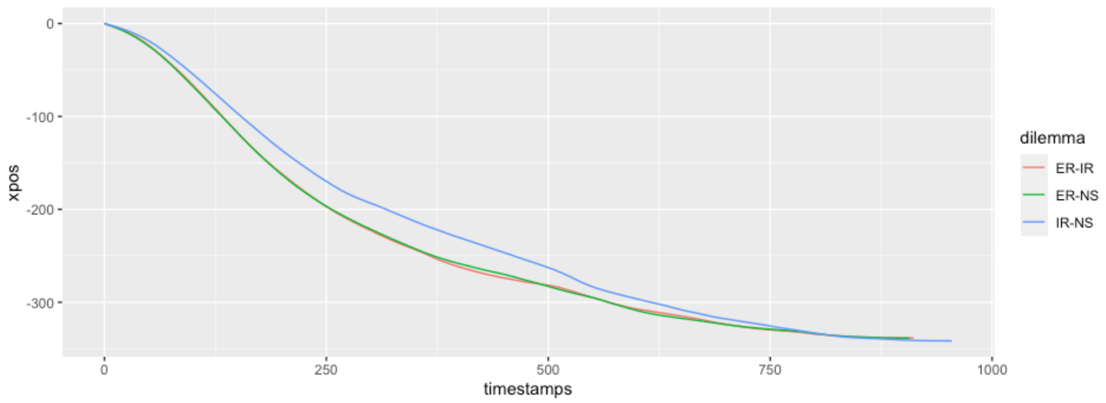

0
-100
-200
-300
xpos
0
250
500
750
1000
timestamps
dilemma
ER-IR
ER-NS
IR-NS

d)

Figure 3.6

Trajectory indices per Dilemma, a) MAD, b) MDabove, c) AUC, d) RTs

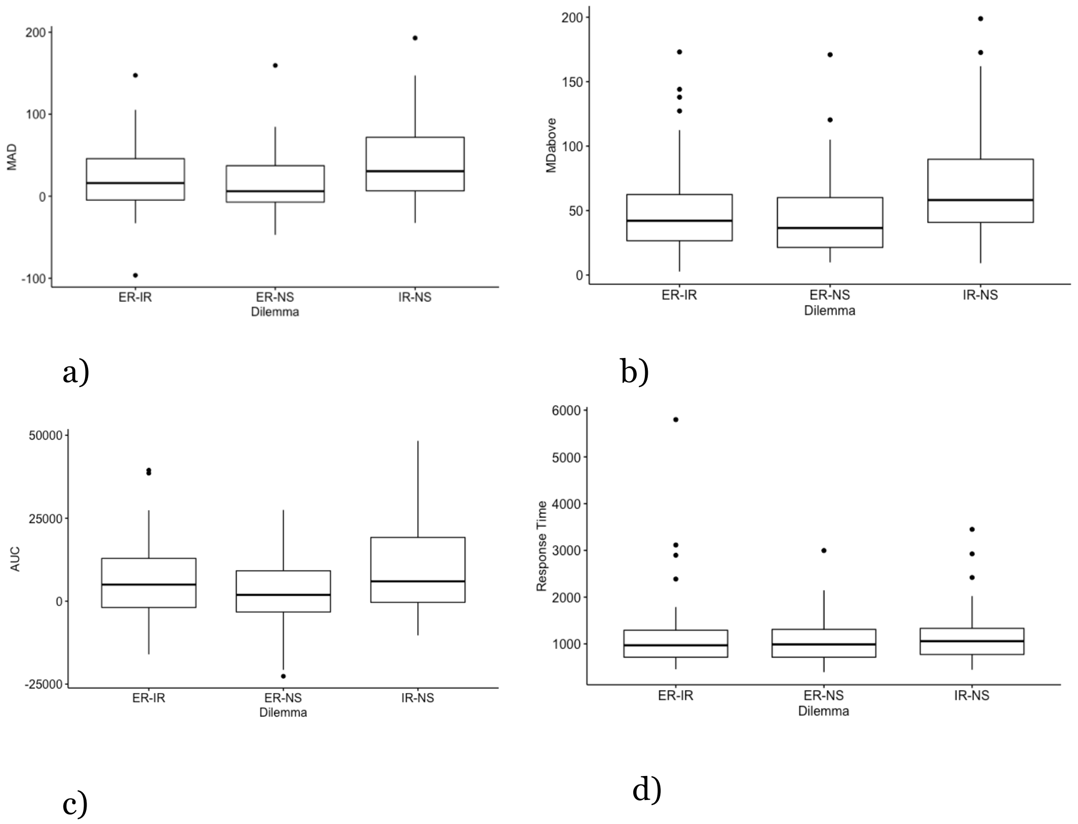

Table 3.3

Frequencies of trajectory types per Dilemma and Age Group.

	Straight	Curved	cCoM	dCoM	dCoM2
ER-IR	805	140	30	24	4
ER-NS	840	136	35	14	5
IR-NS	739	145	44	39	7
Children	616	233	47	27	5
Adolescents	748	105	22	24	1
Adults	1020	83	40	26	10

An illustration of all trajectories clustered in different types can be seen in Figure 3.7. There was no significant DilemmaXAge Group interaction effect on the types of trajectories ($\chi^2(4) = 1.323$, $p = .858$). There was a significant effect of Dilemma ($\chi^2(2) = 14.490$, $p <.001$), and specifically the IR-NS dilemma was significantly different from ER-NS ($p<.001$) and ER-IR ($p = .013$). To look further into these differences, we compared each type of trajectory per type of Dilemma separately. We found that the IR-NS dilemma marginally differed in straight trajectory paths ($F(2, 110) = 2.978$, $p = .055$), especially when compared to ER-IR ($p = .003$). Also, the IR-NS dilemma had significantly more trajectories that showed a discrete change of mind ($F(2, 110) = 2.863$, $p = .004$) and specifically more than the ER-IR ($p = .003$). There was also a significant main effect of Age Group ($\chi^2(2) = 13.261$, $p = .001$). Children differed from adolescents ($p = .02$) and adults ($p < .001$) in their types of trajectories frequencies. Specifically, children differed in terms of their frequency of straight paths ($F(2,54) = 7.874$, $p = .001$), which were significantly less than the adults' one ($p = .001$).

Figure 3.7

All observed trajectories as they have been automatically clustered in five different prototypes by mousetrap, a) Total sample, b) Children, c) Adolescents, d) Adults

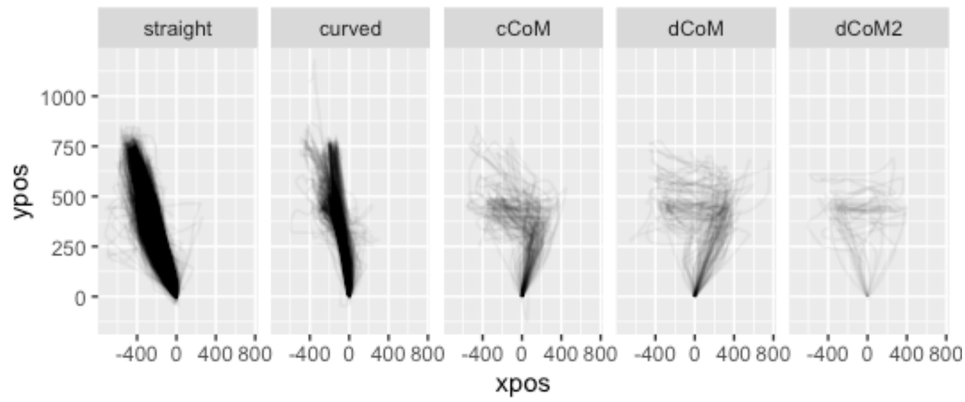

a)

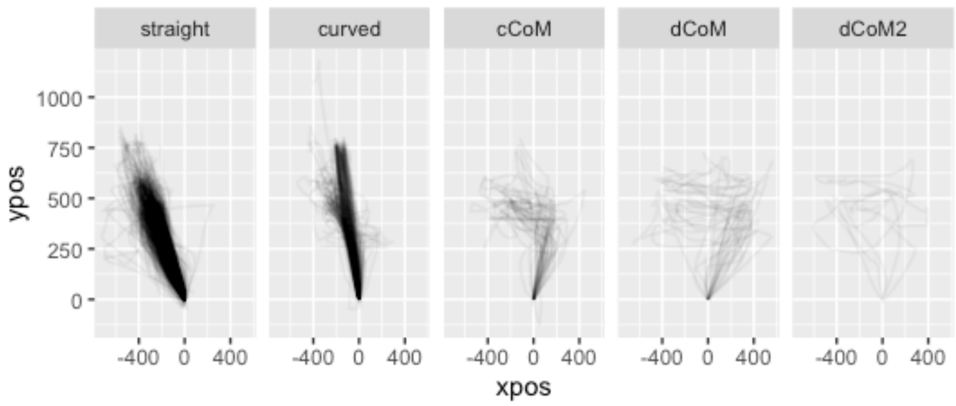

b)

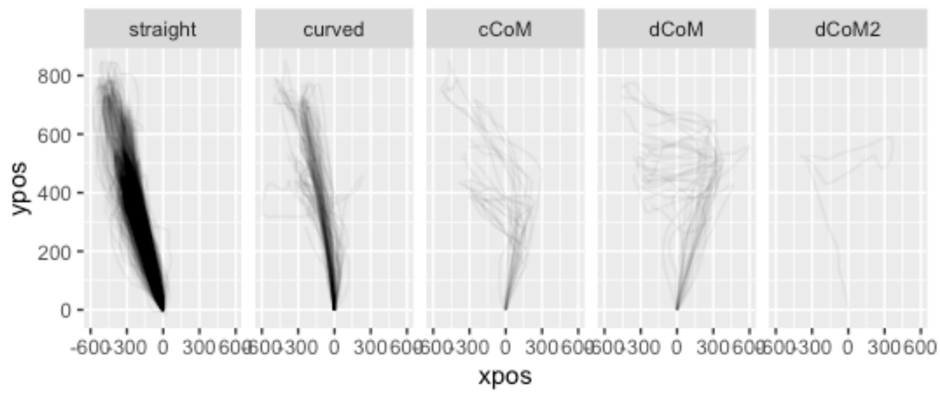

c)

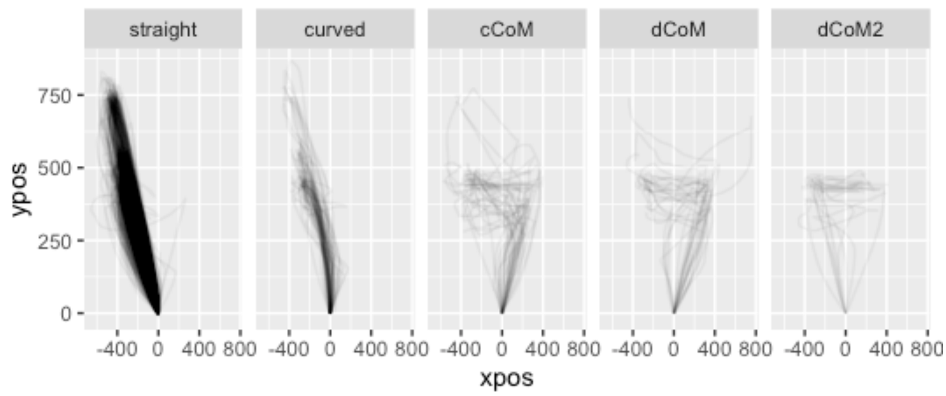

d)

3.2.2.3. *Summary of results*

We observed significant differences in participants' choices overall and per age group. Adolescents chose the ER option more than adults, but not more than children, and the NS option less than adults but not less than children. Children did not differ from the other groups in any of their option choices. Furthermore, children and adolescents chose the ER option more than the IR and the NS options, and only adolescents chose the IR option more than the NS option. Adults, on the other hand, did not differ in terms of their ER and IR choices, and chose ER marginally more often than the NS. When splitting the adolescent group in younger and older participants, we found that the younger group chose the ER more than the IR and NS options (which did not differ), while the older group chose ER more than IR and NS, and IR more than NS.

Regarding movement parameters, the type of dilemma significantly affected all three curvature measures (MAD, MDabove and AUC), RTs and trajectory types. Specifically, in all age groups, participants' trajectories deviated more towards the alternative option in the IR-NS dilemma than in the other two, which did not significantly differ. Their movements in the IR-NS dilemma were also significantly slower compared to the other two dilemmas, and their trajectory paths showed a discrete change of mind more often than in the ER-IR dilemma. Age groups significantly differed in their movement complexity. Specifically, children did more x-flips, and their trajectories followed a straight path significantly less often than adults.

3.2.3. Discussion of Experiment 1

The results suggested that all groups were interested in completing the main goal and achieving the external reward. Furthermore, children and adults were not drawn more to the option with informative value than the option only offering novel stimulation, which was the case only for adolescents. However, we did not find a difference in the frequency of adult choices when they were choosing between the ER and the IR option. The adult approach could be considered the optimal strategy, as the design of the task allowed for achieving both the external reward and the learning goals, with efficient planning of the future choices early in every block. This strategy may therefore indicate mature cognitive control or greater experience present in the adults. At the same time, adolescents showed a different pattern to adults, choosing the ER option more than the IR option. This is consistent with adolescents' previously reported preferences for external rewards (Somerville et al., 2017). However, it is possible that some exploratory strategies were hidden by the fact that the procurement of the external reward led to the termination of the block (especially in adolescents who favoured the ER choice). Therefore, we decided to change this in our next experiment. Furthermore, when we split the adolescents group we found that younger teens were less interested in the IR similar to the one-goal approach that children had, but older teens preferred IR to NS, getting closer to a possibly more mature strategy. However, since the split adolescent groups consisted of only 10 participants each, so it might be that the observed difference was heavily influenced by individual differences and not revealing of a meaningful group developmental difference.

In addition, we observed a clear difference between different dilemmas in several mouse-tracking parameters, suggesting a difference in the decision-making process in the IR-NS dilemma, especially as compared to the ER-IR one. The greater deviation towards the alternative option, as shown in MAD, MDabove and AUC suggests greater difficulty in committing to one of the two exploratory options, a difficulty which particularly materialises in children's equal preference for IR and NS and adults' marginal differences in this dilemma. While these effects were consistent across all age groups, it is possible that they were driven by a few individual participants with more stable dispositions or preferences. Therefore, we aimed to investigate these individual differences more explicitly in Experiment 2.

Moreover, there were interesting differences in the types of trajectories observed in children as compared to the older groups (differences also observed in the elevated number of x-flips in children). These differences in trajectory variability might relate to differences in the integration of values during the decision process at different ages. From a motor control perspective, movements or trajectories that appear continuous on the surface, might hide small sub-movements that can only be revealed by small changes in velocity or acceleration (Dotan, Manyel, & Dehaene, 2018; Wulff et al, 2021). These sub-movements might be more visible when motor control is still developing, and thus materialise in more discrete changes of direction as children integrate conflicting motor plans in their movement.

Lastly, although we did observe effects in the mouse tracking parameters, as it has previously been seen in value-based decision-making tasks (e.g., Koop & Johnson, 2013), no differences were captured in two of the tasks dilemmas (ER-IR, ER-NS). The importance of design choices in mouse-tracking studies has been emphasised repeatedly (e.g., Wirth et al., 2020). Our study had no time limit for the trial onset and, as a result, the decision process might have preceded the onset of movement. Time-limited onset has been previously documented to affect movement parameters (Scherbaum & Kieslich, 2018). Thus, in Experiment 2, we made small design changes, including setting time limits for response, in the hopes of better capturing the decision-making process.

3.3. Experiment 2

In our second experiment, we selected participants with a narrower age range (children 5-7 years old, adolescents 13-15 years old and adults up to 35 years old). We also implemented small design changes in the mouse tracking task (see Procedure).

Based on the observations from Experiment 1, we decided to allow for a longer temporal horizon in the decision-making task; i.e., each experimental block would need more trials to end, and to lead (or not) to the external reward. Longer temporal scales for reward acquisition have been shown previously to favor exploratory behaviour (Sadeghiyeh, Wang, Alberhasky, Kyllo, Shenhav, & Wilson, 2020). We would expect this to be especially reflected in children's and adolescents' choices, which showed strong ER preferences in Experiment 1; they will may make more IR and NS choices as they will have more choices to spare. In terms of their movement parameters, we expected larger conflicts than in Experiment 1 because the exploration options might have become more attractive due to the aforementioned procedural change.

Executive functions

In this experiment, we were also interested in looking more deeply into individual differences between participants. Specifically, we decided to include separate executive functions measures as a means of further uncovering possible connections between cognitive control and exploratory behaviour. In addition, the presence of correlations between executive functions (EF) and specific movements parameters (previously documented; e.g., Erb et al., 2016) would further validate the use of mouse-tracking to approach our experimental question. We were mostly interested in inhibition skills as a causal factor because the control expected by participants to achieve the optimal balance between exploration and exploitation involves choosing the best option in every trial and inhibiting movements driven by the attractive alternative options (either intrinsically or extrinsically rewarding). However, keeping the rewarding goals in mind also involves working memory, and being able to choose based on an overarching goal while navigating different dilemmas possibly also involves flexible response updating. Thus, we included and inhibitory control task, a switching task and a working memory task in our study.

Importantly, EF functions have been measured with a variety of tasks, each focusing on different components and developmental stages. For example, inhibition has been widely measured with Go/No-Go tasks (e.g., Kindlon, Mezzacappa, & Earls, 1995), where participants are trained in a dominant response and then either their ability to inhibit this response is measured, or their ability to readily respond correctly (e.g, performance on omission vs. commission errors has been found to correlate with different control abilities; Bezdjian, Baker, Lozano, & Raine, 2009). Other inhibition tasks are based on pre-existing dominant responses, for example Stroop-like inhibition tasks (Ikeda, Okuzumi, & Kokubun, 2014) or the Flanker task (Zelazo, Anderson, Richler, Wallner-Allen, Beaumont, & Weintraub, 2013). In our study, we will use a version of the Go/No-Go task, measuring the participants errors of commission, which have been previously associated with impulsivity (Bezdjian et al., 2009). Similarly, task switching has been assessed in various ways, but most of these are a variation of the Dimensional Change Card Sorting task (DCCS; Zelazo, 2006), which expects participants to flexibly change their learned responses based on a rule change across blocks of trials. In our study, we designed a version of the DCCS task using different cars and boats of different colours, broadly following the design of the equivalent task by Howard and Melhuish (2017) for their Early Years Toolbox (YET). The stimuli presented remained the same but expected responses changed based on the task rule. We also relied in on the YET for the design our working memory task. This task included auditory instructions and the goal was to choose the correct shape among many. The instructions had growing difficulty, which participants should remember to be able to choose the correct shape. We expected participants with larger choice conflicts and more unbalanced choices (e.g., more NS choices) to score worse in general on the EF tasks, especially the Go/No-Go task.

Individual traits

Curious behaviour (i.e., increased exploration) has also been associated with differences in people’s personalities. These associations have predominantly focused on stable dispositions rather than state-dependent exploration or developmental differences, which were our main focus in our first experiment. For example, a series of studies by Litman and colleagues have associated different types of curiosity with sensation seeking, anxiety and ambiguity tolerance (Litman, 2010; Litman, Collins, & Spielberger, 2005), while more recently aspects of trait curiosity have been decomposed by Kashdan et al. (2018). In this second experiment, we aimed to discover whether two specific and supposedly stable dispositions can predict participants' exploratory strategies; specifically, (i) complexity preference, and (ii) unpredictability preference.

Instead of using questionnaires, we designed two short tasks in which participants could choose how long they would explore options of different visual complexity and predictability. Visual complexity was chosen because of the seminal work suggesting that people preferably attend to images of intermediate complexity as compared to high or low complexity (e.g., Berlyne, 1958; Kidd et al., 2014). We expected participants' preferences for intermediate visual complexity to correlate with more directed exploration (IR choices) in all ages. Stimulus unpredictability (i.e., uncertainty in the temporal domain) is the most common variable manipulated in "armed bandit" experiments. It directly affects learning and exploration by making environments more or less easily learnable (e.g., Poli et al., 2022; Blanco & Sloutsky, 2021).

In our task participants could choose to explore choices at three levels of predictability (100%, 70% and 50% probability of the same stimulus appearing), all of which were learnable but at different rates. We expected age differences, such that children might explore the 50% option more, whereas adolescents and adults would prefer the 70% option. These preferences for each age group were expected to correlate positively with more IR choices for each participant.

3.3.1. Methods

3.3.1.1. Participants

In this experiment, the participants were 22 children (9 females, mean age: 6.09 years), 20 adolescents (10 females. mean age: 14.47 years) and 24 adults (11 females, mean age: 32.5 years). All participants were neurotypical and had normal or corrected-to-normal vision. The participants were volunteers recruited mainly from word-of-mouth and through the Birkbeck Babylab database. They received a £5 Amazon voucher for their participation.

3.3.1.2. Mouse-tracking task

3.3.1.2.1. Design

Participants of all age groups had to complete four experimental blocks of 27 trials each. In each trial, two of the total three options (selected randomly) appeared on the screen. The participants had to choose an option (in a *2 alternative forced choice* format) in one of the following dilemmas: External Reward VS Informational Reward (ER-IR), External Reward VS Novel Stimulation (ER-NS), or Informational Reward VS Novel Stimulation (IR-NS). If the External Reward (ER) option is chosen for a total of nine times, the participant won the reward. The trials continued until a total of 27 trials was completed. This means the maximum choices for each option could be 18.

3.3.1.2.2. Stimuli and procedure

The stimuli and procedure were almost identical to those in Experiment 1, with the following important changes, applied both on the training and the test phase: participants had a time-limit of 6s (children) or 2s (adolescents and adults) to start moving the bone towards the (dog) choice options (i.e., to start the trial). The two options did not appear unless the bone had been moved a minimum of 4 pixels above its initial position. If the time-limit was exceeded, participants were informed that they were too slow and the trial started over again. An equivalent time limit was set for the answer – participants had to pick an option in 6s (children) or 2s (adolescents and adults). They received the same warning if they were too slow with their answer. No punishment was inflicted in terms of remaining lives for the completion of the block.

3.3.1.3. Secondary tasks

All secondary tasks in our experiments were created and hosted using the Gorilla Experiment Builder (www.gorilla.sc; Anwyl-Irvine, Massonnié, Flitton, Kirkham & Evershed, 2018).

3.3.1.3.1. Inhibition task

We used a child-friendly Go/No-Go task adapted from Bezdjian et al. (2009). Specifically, we replaced the letters which served as Go and No-Go stimuli in the study by Bezdjian et al. (capital Ps and capital Rs respectively) with animal images (a hedgehog as a Go stimulus and a bear as a No-Go stimulus) and the array of stars with an array of flowers. The number of trials and breaks was also adapted to make the total experimental time efficient for testing children online: while experimental blocks lasted 160 trials in the original study, they only lasted 40 trials in our study. The Go/No-Go ratio was kept the same. The same version of the task was used for all age groups.

Design

Participants had to complete three experimental blocks of 40 trials each. The ratio of Go and No-Go trials was 80% Go – 20% No-Go trials. The Go/No-Go trials were pseudorandomized, making sure no subsequent No-Go trials were presented.

Stimuli and procedure

Participants were instructed to try and catch the mole which appeared among three flowers, in one of four random positions in a square grid (Go Stimulus; Figure 3.8a), but they had to avoid catching the bear (No-Go stimulus; Figure 3.8b). They had to press the Space key as soon as the mole appeared on the screen but not when a bear appeared. The experiment started with 10 practice trials, which were repeated if the participant was too slow or if they made an error at both of the No-Go trials (which were pseudorandomised to appear after five Go trials to create a prepotent response). Each trial started with a presentation of a block of four flowers for 1500ms (Figure 3.8c), followed by the target stimuli for 500ms. The response screen showed four flowers again for 1200ms. Participants had to respond during this time limit, or the trial was recorded as an error. After each response, participants received feedback (as a thumbs up or a thumbs down image) for 200ms, both in practice and experimental trials.

Figure 3.8

Go/No-Go task stimuli, a) Go stimulus, b) No-Go stimulus, c) Masking Image

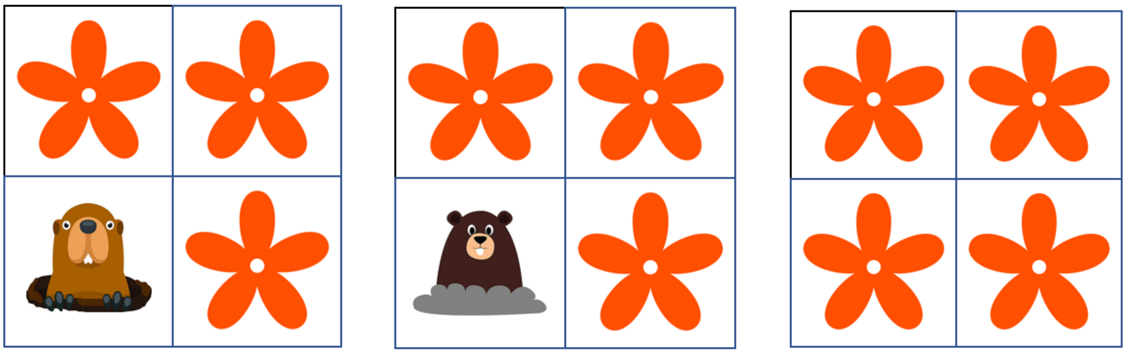

3.3.1.3.2. Switching task

This task was adapted from the YET toolbox's 'Rabbits and Boats' task (Howard and Melhuish, 2017). First, we used a simpler interface. Secondly, the participants in our task had to decide between cars and boats (Vehicle rule), or blue and red (Colour rule). Furthermore, the change of rule was communicated through a shape before the beginning of each block of trials (see Procedure below). The total number of blocks was also decreased as compared to the original.

Design

Participants had to complete four experimental blocks, two with each rule (Vehicle and Colour). Each block had six trials. The task rule sequence was always Vehicle-Colour-Vehicle-Colour. Accuracy was measured as the number of errors at the incongruent trials following a rule change. A trial was considered incongruent when pressing the response that was expected was the opposite from the corresponding image at the top of the screen (see Table 3.4).

Table 3.4

Switching task design according to stable on-screen categories

	Vehicle **Type of Response**	**Colour** **Type of Response**
Red car	Incongruent	Congruent
Blue car	Congruent	Incongruent
Red boat	Congruent	Incongruent
Blue boat	Incongruent	Congruent

Stimuli and procedure

Participants started the task with two practice blocks of three trials per task rule. At the beginning of all trials, they watched a blue car in the top left of their screens and a red boat in the top right (Figure 3.9a). Then the shape signaling the rule was presented in the middle of the screen for 400ms. A yellow triangle was used to signal the Vehicle rule, whereas a green circle was used to signal the Colour rule (Figures 3.9b and 3.9c). A verbal reminder was also given along with the presentation of the shapes. When the shape disappeared, the target stimulus appeared at the middle bottom of the screen and remained until a response was given (Figures 3.9d and 3.9e). Participants had to respond by pressing the Z or M keys on their keyboards to categorise the stimuli. Feedback was given during the practice trials but not during the experimental trials.

Figure 3.9

Switching task stimuli, a) Background categories screen, b) Rule 1 (Vehicle type), c) Rule 2 (Colour), d) Correct categorization of blue boat stimulus based on Rule 1, e) Correct categorization of red car stimulus on Rule two.

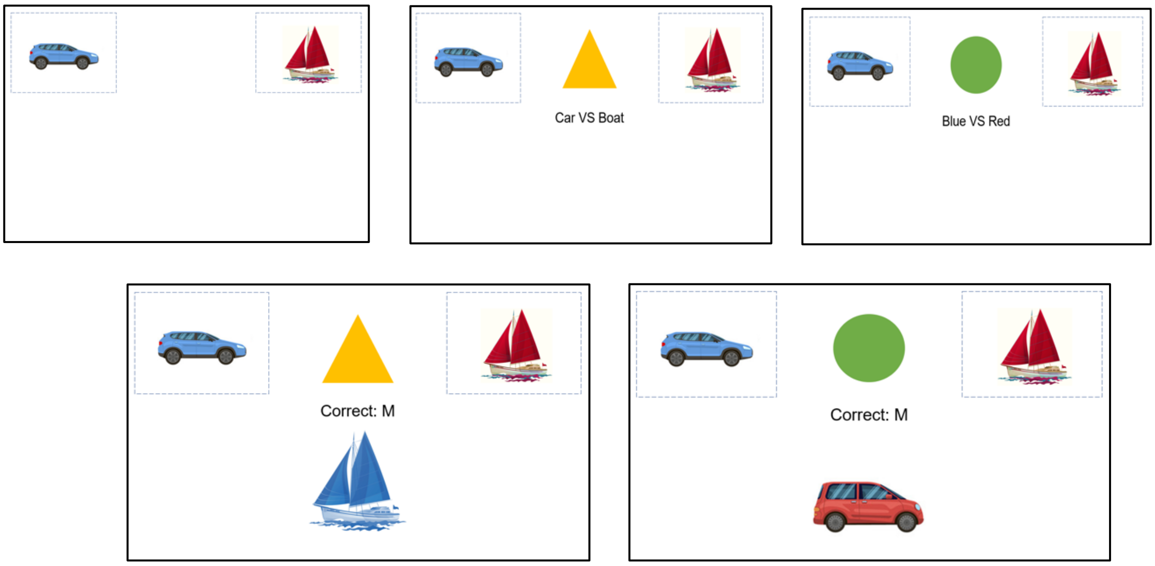

3.3.1.3.3. Working memory task

The task was similar to a Direction Following task (Im-Bolter, Johnson, & Pascual-Leone, 2006), measuring specifically phonological memory, as it required participants to keep in mind features of shapes (shape, size and colour) which were described to them verbally. They were then asked to choose objects which did *not* have the aforementioned characteristics. As the task progressed, the features increased in number, starting from simple descriptions, e.g., 'Find a shape that is not red', to much more complicated ones, e.g., 'Find a shape that is not big, not red, not a triangle and not a circle'. The task is an adapted version of the 'Not This' task of the YET (Howard & Melhuish, 2017), with the main change being the number of trials per block (our task had 3 trials), to make the task shorter for children and prevent fatigue.

Design

The task consisted of a maximum 7 blocks/levels of complexity, consisting of three trials each. Each level involved adding an extra feature in the description of the shapes. However, the total duration varied based on a participant's performance. If a participant made two mistakes on the same level or three mistakes in total, the task was terminated. Performance was measured as the maximum level they managed to complete, plus 0.33 for every correct trial after that before they lost (e.g., they scored 6.33 if they only got one trial correct on level 6 and missed the next two).

Stimuli and procedure

Participants started the task with three practice trials (two Level-1 and one Level-2 example). On each trial, participants had to press a button presented at the centre of the screen in order to hear the recorded description of the shape. They could only hear the recording once. After the recording, the screen remained blank for 3000ms and it was followed by a screen where 20 shapes were presented in button-like form (Figure 3.10). The participants had to click on the shape that did not have any of the features described at the recording. No time limit was set for them to answer.

Figure 3.10

Working memory task shapes configuration.

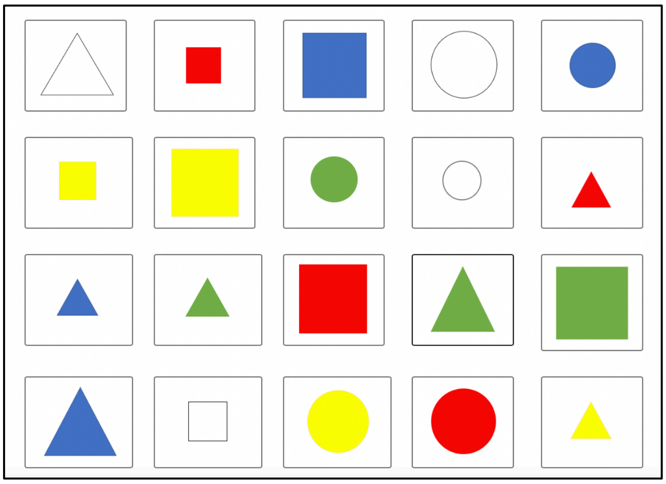

3.3.1.3.4. Complexity preference task

Drawing from older stimulus preference tasks, measuring participants' looking time or explicit preference/liking judgements on complexity (e.g., Day, 1966; Munsinger, Kessen & Kessen, 1964), we designed a task in which participants could choose to view/explore for more or less time four abstract stimuli of growing complexity. There was a minimum time of exploration per trial, in order to make participants' preferences more obvious.

Design

The task consisted of 12 trials. During each trial, participants could choose between two alternatives of different complexity (2AFC task). We used stimuli of 4 levels of complexity (see *Stimuli and Procedure*) and no image was presented twice. This meant that each level was presented 6 times, in 3 different dilemmas, and contrasted with levels of greater or smaller complexity. Preference was measured as the total frequency of choices for each level.

Stimuli and Procedure

We used a number of symmetric stimuli produced by Gartus and Leder (2017). We clustered these into four complexity levels, based on the human and a computational model's ratings provided in Gartus and Leder (2017). The stimuli were coloured to be more appealing to children (the colours were randomized for each level). Example stimuli for each level is shown in Figure 3.11a.

At the beginning of the task, participants completed a practice trial during which they were encouraged to explore the options they liked more. On each trial, participants were initially presented with a configuration of two closed windows (left-right; Figure 3.11b). The windows remained closed for 3000ms, then both opened for 1500ms, revealing the abstract shapes of the trial (Figure 3.11c). Then they closed again, and the participants had to choose which window they wished to open by pressing the left or right arrow on their keyboards. When they pressed the key, the window of their choice opened for 1000ms, revealing the image, and then closed again. Then a new choice had to be made. Participants had to keep choosing for a minimum of 6 times. Finally, a button appeared in the bottom right of the screen allowing them to proceed to the next trial. However, if they wanted to explore the current stimuli more, they could keep choosing one of the windows for up to 6 more times (i.e., 12 times in total).

Figure 3.11

a) Example of experimental stimuli at four different complexity levels (adapted from Gartus and Leder, 2017), b) Initial screen, hiding the experimental stimuli, c) Presentation of both stimuli at the beginning of each block, d) Revealed stimulus after choice has been made.

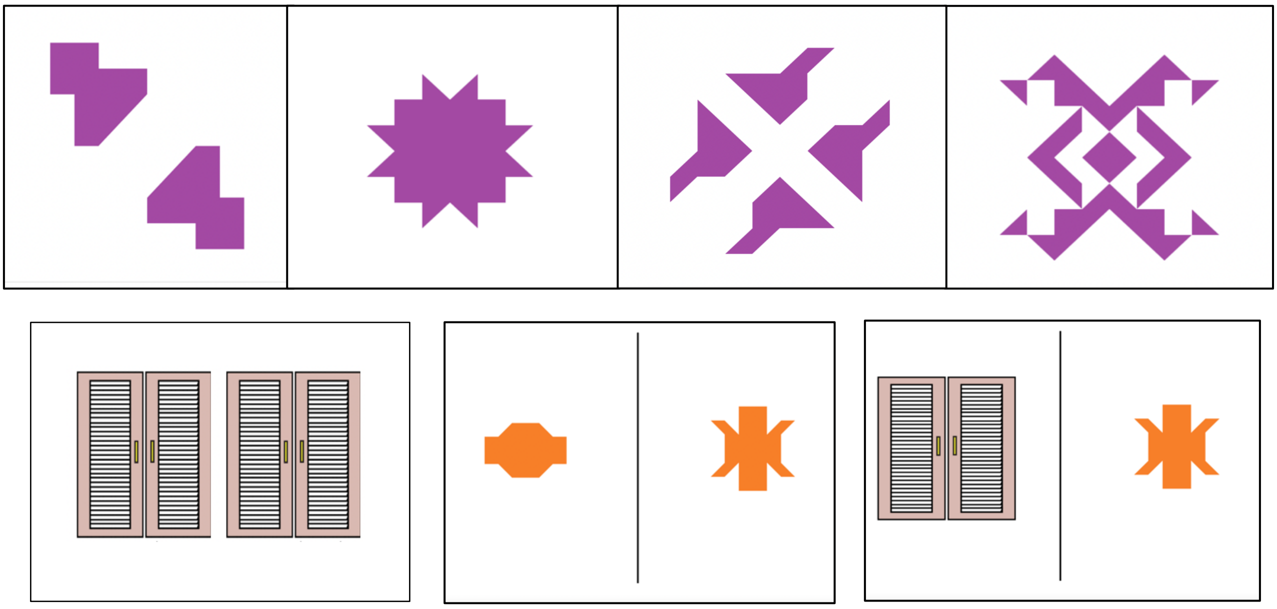

3.3.1.3.5. Unpredictability preference task

As with the complexity preference task, we chose to design this task as a decision task between options that were more or less predictable (i.e., they allowed participants to learn the statistics of stimulus presentation more quickly or more slowly). Three different options were presented simultaneously: (i) one which always included the presentation of the same abstract image every time that it was selected (100% predictable), (ii) one in which the same stimulus was presented on 70% of the trials and another (but also stable) stimulus was presented on the remaining 30%, and (iii) one in which two stable stimuli were presented for 50% of the trials each.

Design

Participants had to complete three blocks of 20 trials each. During each trial, participants could choose between three options. They had to keep choosing until the block ended. Preference was measured as the total frequency of choices for each level of predictability.

Stimuli and procedure

For this task we used a subset of the stimuli used in the complexity preference task; specifically, we only used stimuli of the same level of complexity for each exploration block (e.g., five level-4 stimuli for Block 1, etc.; see Figure 3.12a).

As the task had a similar layout to the complexity one, no practice trials were used. Participants were instructed to choose the option that they wanted to see by pressing the keys 1, 2 or 3 on their keyboards. On each trial, they were initially presented with a configuration of three closed windows (Figure 3.12b). The windows remained closed until the participants pressed the key associated with their preference. The stimulus then appeared for 1000ms and was hidden again. The participant was expected to make choices until 20 trials were completed, and then a button appeared in the bottom right of the screen allowing them to proceed to the next block.

Figure 3.12

a) Set of stimuli used in the same block of the unpredictability task, b) Hiding windows configuration and predictability, c) A stimulus shown after window 2 has been chosen

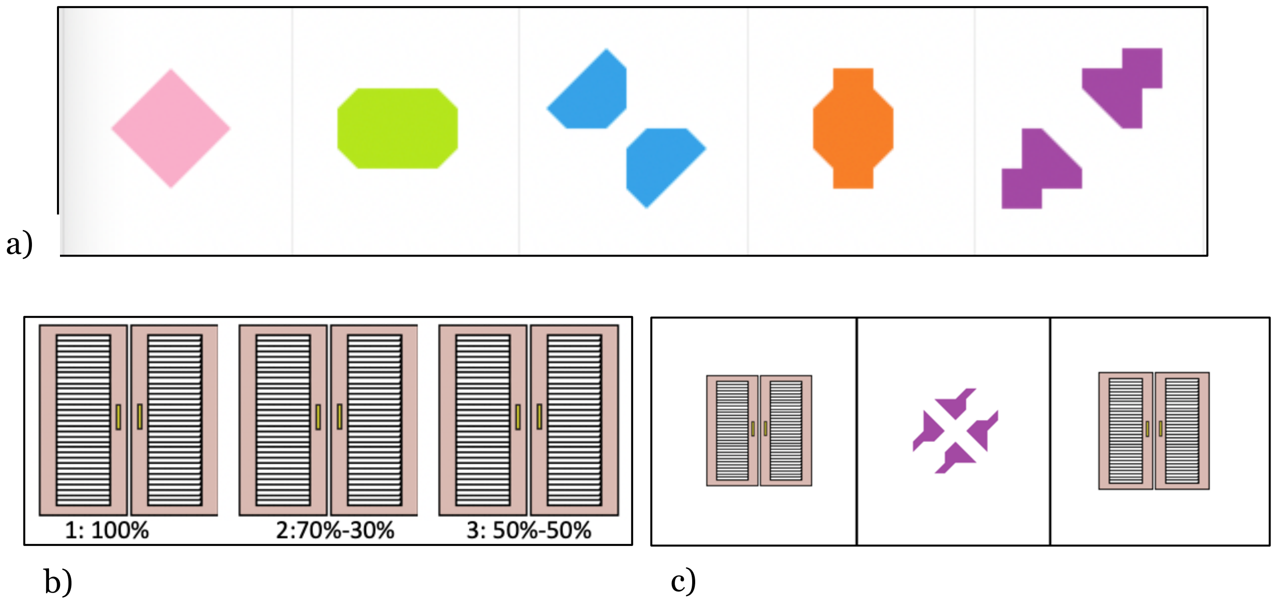

3.3.2. Results

No participants were excluded from this study; however, one child did not complete the final block of trials in the mouse-tracking study due to technical issues with their computer. Outliers in the separate tasks were excluded from the relevant analyses. The exclusion criteria will be discussed separately below for each task.

3.3.2.1. Choices

Figure 3.13 and table 3.5 show the mean number and SDs of choices of each of the three options for each age group. We first analysed the choices each age group made and found a significant interaction of Age Group and Option: $F(4, 122) = 10.853$, $p<.01$. Post-hoc Bonferroni-corrected comparisons showed that in total, participants chose the ER option more than the IR option ($p=.038$) and the IR option more than the NS option ($p<.001$). Simple effects comparisons showed that children chose the NS option more than adolescents ($p = .004$) and adults ($p<.001$), whereas groups did not significantly differ in their other choices. Analyses within age groups showed no differences in children's choices: children chose the ER option equally to the IR option ($p = .092$) and the NS option ($p = .694$). Adolescents chose ER significantly more than NS ($p < .001$) and IR significantly more than NS ($p = .003$), while their ER and IR choices did not differ ($p=.886$). Finally, adults also chose ER significantly more than NS ($p<.001$), and IR significantly more than NS ($p<.001$) but their ER and IR choices did not differ ($p = .789$).

Table 3.5

Means and SDs of choices per Option for each age group

	Children	*Adolescents*	*Adults*
External Reward	9.716 (2.106)	10.342 (1.822)	10.281 (2.310)
Informational Reward	8.602 (2.269)	9.829 (1.868)	10.120 (3.300)
Novel Stimulation	8.989 (2.374)	6.474 (2.195)	5.837 (2.500)

Figure 3.13

Number of choices per Option for each age group in Curiosity task

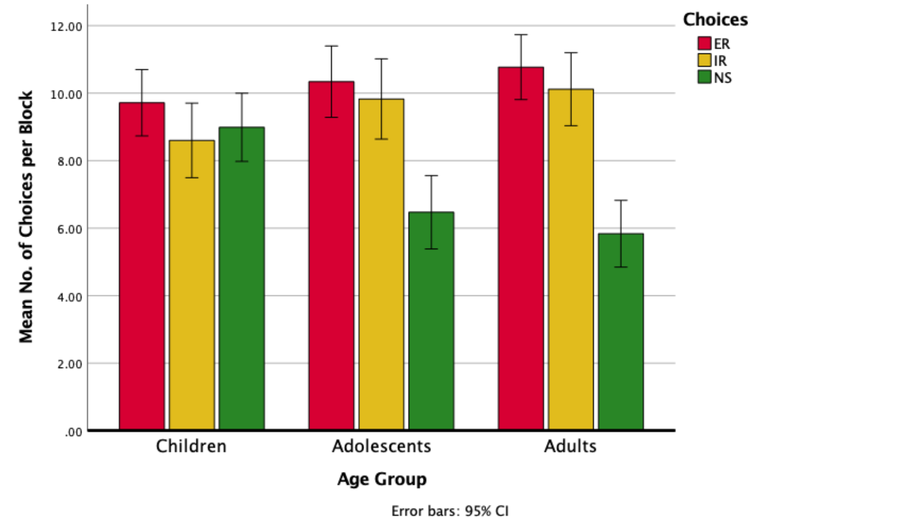

To examine if participants were affected by the longer decision horizon, we also analyzed their choices up to the completion of the tower (i.e., the ER goal prerequisite), if they had chosen to complete it, or one of the other two choices (i.e., if they completed the puzzle, or chose the NS for nine times). This way, the behaviour could be comparable to what participants did in Experiment 1, when they knew that they could not keep exploring if they acquired the ER. We analysed the choices each age group made and found a significant interaction of Age Group and Option: $F(4, 124) = 4.665$, $p = .002$. The main effect of Option was also significant: $F(2,124) = 12.945$, $p<.001$. Post-hoc Bonferroni-corrected comparisons showed that in total, participants chose the ER option more than the IR option ($p=.003$) and the NS option, while the IR and the NS options did not differ ($p = .158$). Overall, children chose the IR option significantly less than adults ($p = .0.48$) and the NS significantly more than adults ($p = .013$). We further examined the effect of option on each age group. Children's choices did not differ, whereas adolescents chose ER significantly more than IR ($p =.011$) and NS ($p = .035$) but their IR and NS choices did not differ ($p = .789$). Adults chose ER significantly more than NS ($p<.001$), and IR significantly more than NS ($p<.001$) but their ER and IR choices did not differ ($p = .819$).

3.3.2.2. Mouse-tracking data analyses

Here, we followed the same procedure as in Experiment 1 to extract features from the raw mouse-tracking data. Table 3.6 shows means and SDs for the five indices we chose to compare between each dilemma.

Table 3.6

Means and SDs for movement parameters per Dilemma and Age Group

		Children		**Adolescents**		**Adults**	
Features[a]		Mean	SD	Mean	SD	Mean	SD
Maximum Absolute Deviation (MAD)	ER-IR	133.60	191.62	167.62	199.80	180.10	197.63
	ER-NS	135.47	200.70	188.94	215.41	186.14	195.72
	IR-NS	142.46	204.23	169.25	207.12	184.64	197.98
Maximum Deviation above ideal line (MDabove)	ER-IR	161.20	159.47	180.68	184.11	187.64	189.12
	ER-NS	167.42	163.32	202.79	198.84	193.43	187.43
	IR-NS	173.30	171.20	184.19	191.21	192.55	189.24
Area Under the Curve (AUC)	ER-IR	39358.80	60626.44	48670.43	68616.82	56187.19	74058.19
	ER-NS	37838.06	59274.46	53952.11	72629.98	58374.96	74034.01
	IR-NS	38907.45	62177.45	48398.64	71127.28	58746.30	76899.29
x-flips	ER-IR	2.10	1.84	1.29	1.22	1.25	1.10
	ER-NS	2.02	1.82	1.44	1.28	1.31	1.20
	IR-NS	2.16	1.96	1.34	1.28	1.37	1.21
Response Times	ER-IR	2590.18	2493.74	810.85	304.33	896.55	342.89
	ER-NS	2541.43	2516.43	842.33	326.51	925.16	442.91
	IR-NS	2597.26	2826.00	823.84	311.86	910.62	343.47

[a]All time related values are presented in milliseconds (ms), all position related values are presented in pixels (px), area (AUC) is displayed in px^2.

There was no significant DilemmaXAge Group interaction effect on MAD ($\chi^2(4) = 3.984$, $p = .408$). Type of Dilemma did not affect MAD ($\chi^2(2) = 2.507$, $p = .285$), neither did Age Group ($\chi^2(2) = 2.930$, $p = .231$). No significant interaction between Dilemma and Age Group was observed for MDabove ($\chi^2(4) = 1.815$, $p = .770$). Type of Dilemma did not affect MDabove ($\chi^2(2) = 2.255$, $p = .324$, neither did Age Group ($\chi^2(2) = 1.902$, $p = .386$). Furthermore, there was no significant DilemmaXAge Group interaction effect on AUC ($\chi^2(4) = 3.444$, $p = .486$). Type of Dilemma did not have a significant effect on AUC ($\chi^2(2) = 1.279$, $p = .528$), neither did Age Group ($\chi^2(2) = 2.457$, $p = .293$. There was a significant DilemmaXAge Group interaction effect on xflips ($\chi^2(4) = 10.525$, $p = .032$) (Figure 3.14). Post-hoc comparisons showed that children did significantly more flips than adolescents ($p<.001$) and adults ($p<.001$). No main effect of Dilemma was observed ($\chi^2(2) = 1.068$, $p = .586$). Lastly, no significant DilemmaXAge Group interaction effect was observed on RTs ($\chi^2(4) = 2.036$, $p = .729$). There was no main effect of Dilemma on RTs ($\chi^2(2) = 1.165$, $p = .559$). There was a significant effect of Age Group on RTs ($\chi^2(2) = 85.602$, $p <.001$). Specifically, children were slower than adolescents ($p<.001$) and adults ($p<.001$), but they were also given more time to respond (6s compared to 2s)[3].

Figure 3.14

Mean number of x-flips per Age Group across all dilemmas.

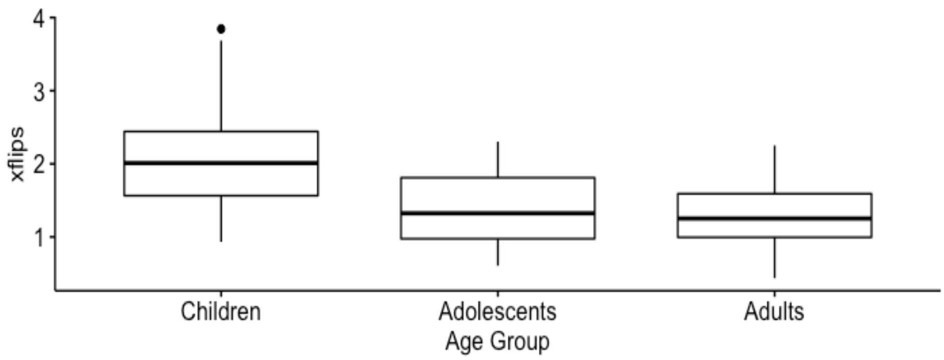

Figure 3.15 shows the automatically clustered trajectories for the full sample, and for each age group. There was no significant DilemmaXAge Group interaction effect on the types of trajectories ($\chi^2(4) = 1.752$, $p = .781$). There was a significant

3 When RTs were standardised, no main effect of Age Group was found ($\chi^2(2) = 0.058$, $p = .971$.

effect of Dilemma ($\chi^2(2) = 15.102$, $p < .001$), and specifically the IR-NS dilemma was significantly different from ER-NS ($p<.001$) and ER-IR ($p = .006$). To look further into these differences, we compared each type of trajectory per type of Dilemma separately. We found that the IR-NS dilemma significantly differed in frequency of straight trajectory paths ($F(2, 110) = 17.027$, $p < .001$), and specifically when compared to ER-NS ($p = .002$). Also, the IR-NS dilemma had significantly more trajectories that showed a discrete change of mind ($F(2, 110) = 4.850$, $p = .011$) and specifically more than in the ER-NS dilemma ($p = .013$). There was also a significant main effect of Age Group ($\chi^2(2) = 13.079$, $p = .001$). Children differed from adolescents ($p = .02$) and adults ($p < .001$) in their types of trajectories frequencies. Specifically, children differed in terms of their straight path frequencies ($F(2,53) = 6.135$, $p = .004$), which were significantly more than the adults' ones ($p = .003$). The difference in curved path frequencies was marginal ($F(2,53) = 2.973$, $p = .060$).

Table 3.7

Frequencies of trajectory types per Dilemma and Age Group.

	Straight	**Curved**	**cCoM**	**dCoM**	**dCoM2**
ER-IR	875	674	311	185	35
ER-NS	842	623	370	187	32
IR-NS	888	616	312	169	34
Children	969	580	290	116	23
Adolescents	752	617	321	168	45
Adults	884	716	382	257	33

Figure 3.15

All observed trajectories as they have been automatically clustered in five different prototypes by mousetrap, a) Total sample, b) Children, c) Adolescents, d) Adults

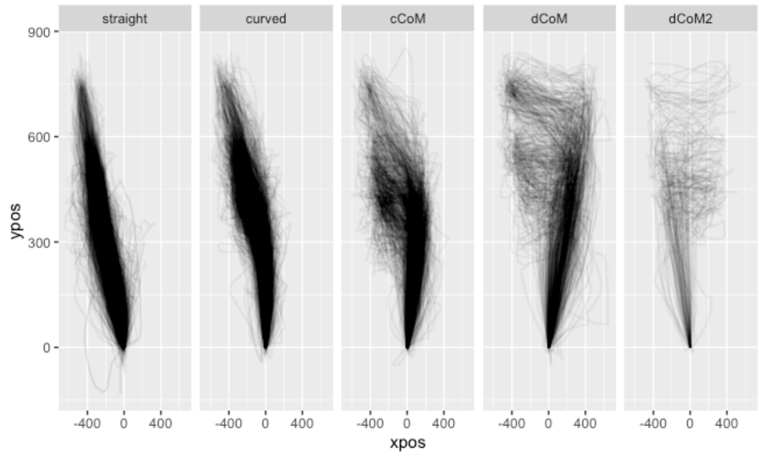

a)

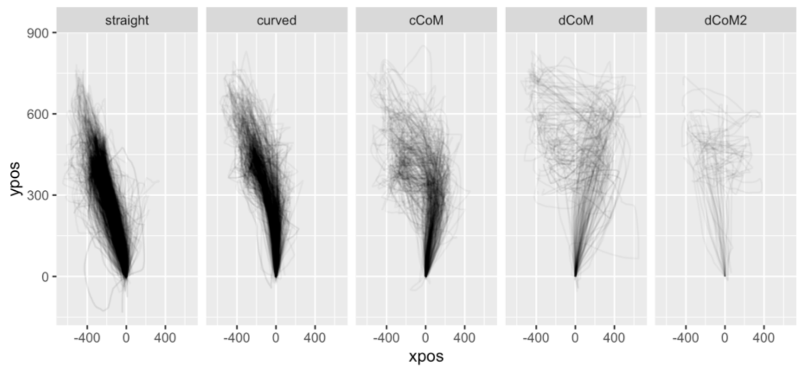

b)

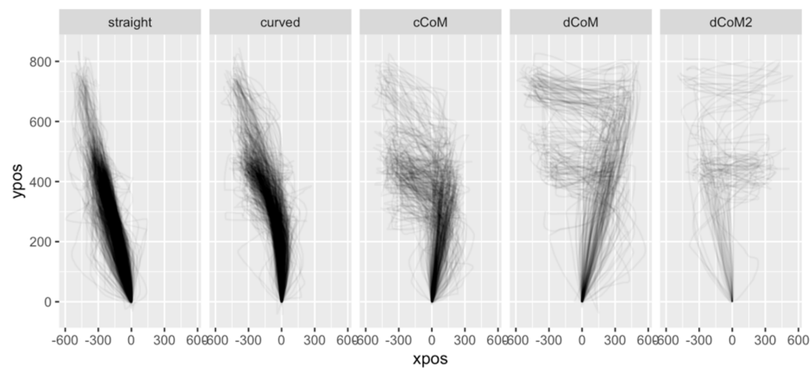

c)

d) 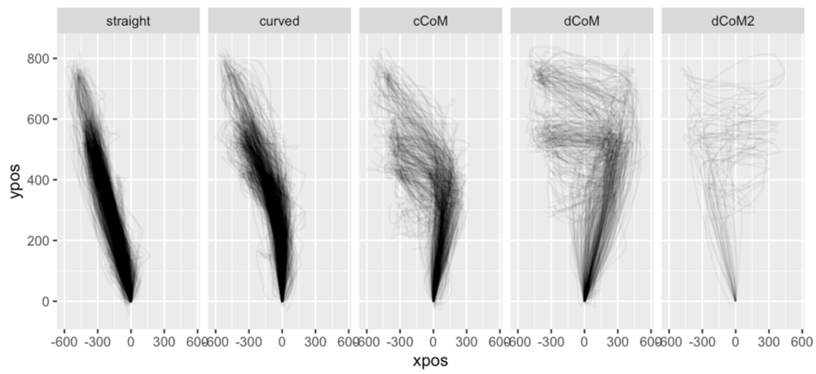

3.3.2.3. Secondary tasks

A subgroup of participants did not complete these tasks due to fatigue (11 children) or connectivity/PC issues (6 adolescents and 4 adults). The final sample for these measures consisted of 13 children, 14 adolescents and 20 adults. No outliers were excluded.

3.3.2.3.1. Complexity preference task

The complexity preference score was calculated as the total number of choices for each level of complexity (1 to 4) across all trials for each participant.

We performed a two-way ANOVA with Age and Level of Complexity as factors (Figure 3.16). There was no significant interaction ($F(6,132) = 1.958$, $p = .076$, but there was a main effect of Level of Complexity ($F(3, 132) = 25.299$, $p <.001$. Specifically, participants chose the less complex option significantly less than the others ($p<.001$). Age did not have a significant effect ($F(2,44) = 2.783$, $p = .076$).

3.3.2.3.2. Unpredictability preference task

The unpredictability preference score was calculated as the total number of choices for each level of predictability (1 to 3) across all trials by each participant.

We performed a two-way ANOVA with Age and Level of Predictability as factors with the number of choices as the dependent variable (Figure 3.17). There was a significant interaction between Age and Level of Predictability ($F(4,80) = 5.082$, $p = .001$). We then analysed age groups separately. Only adults were significantly affected by the level of predictability (F(2,32) = 5.595, p = .008); they chose the less uncertain option less than the more uncertain (p = .016).

Figure 3.16

Complexity preference per Age Group

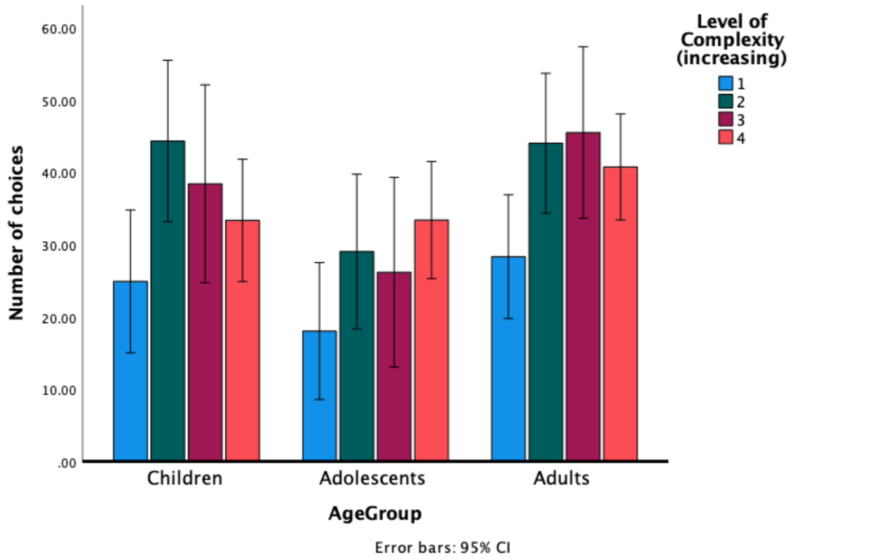

Figure 3.17

Unpredictability choice preference per Age Group

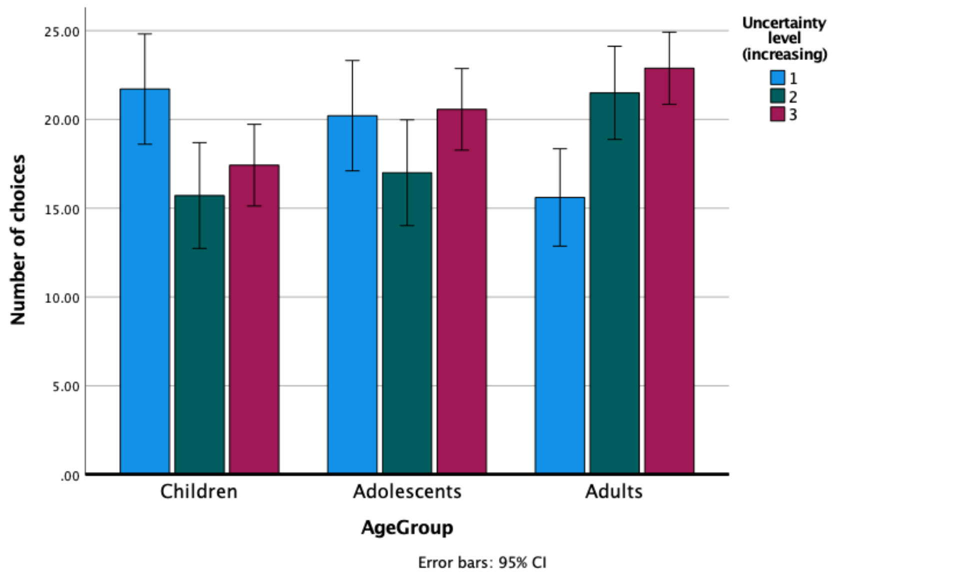

3.3.2.3.3. Go-No Go (Inhibition) task

Task performance was measured as the number of errors in the No Go trials divided by the total number of No Go trials. We performed a one-way ANOVA with Age Group as factor (Figure 3.18a). There was a main effect of Age Group ($F(2,40) = 7.990$, $p = .001$). Specifically, children performed significantly worse than adolescents and adults ($p = .001$ and $p = .002$). However, there was essentially a ceiling effect for the two older groups.

3.3.2.3.4. Switching task

Task performance was measured as the number of errors in Incongruent trials after a set-rule change. We performed a one-way ANOVA with Age Group as factor (Figure 3.18b). There was no significant difference between the groups' performance ($F(2,45) = 0.718$, $p = .494$.

3.3.2.3.5. Memory task

Task performance was measured in the following way: Participants earned 1/3 of a point for every correct answer. They then proceeded to the next level if they got the 2/3 trials of a level correct. The score was the last one accomplished before a total of three mistakes and two consecutive mistakes. For example, a score of 5.33 means the participant completed five full levels and only got one correct trial of the sixth level, followed by wrong ones. We performed a one-way ANOVA with Age Group as factor (Figure 3.18c). There was a significant main effect of Age ($F(2,39) = 14.149$, $p<.001$), with children performing significantly worse than the other two groups ($p = .012$ and $p<.001$), who did not differ from each other.

Figure 3.18

Performance in Executive functions tasks per Age Group. a) Inhibition task, b) Switching task, c) Working Memory task

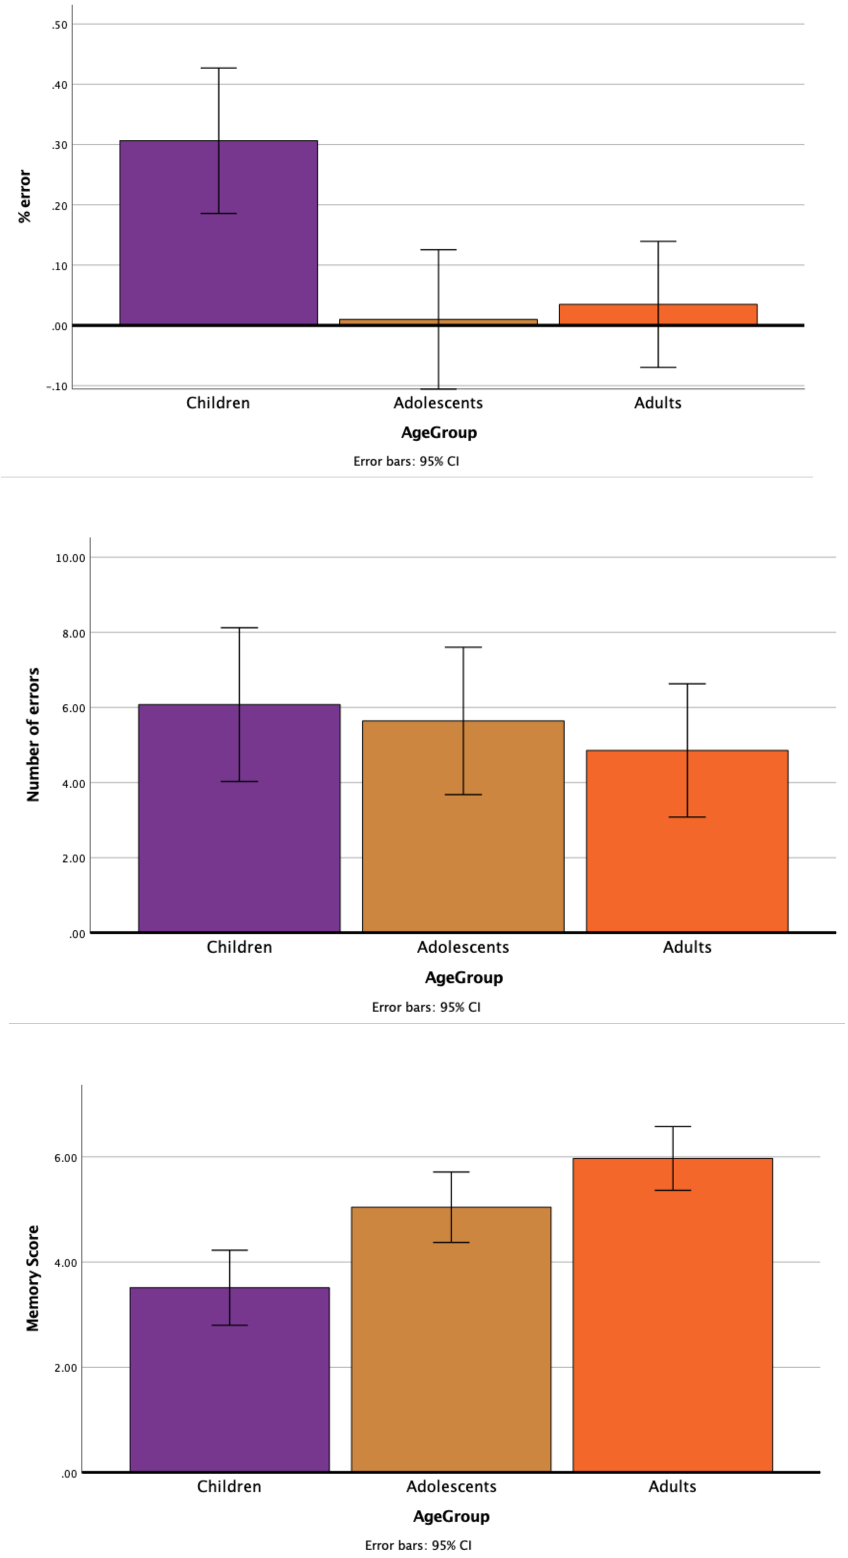

3.3.2.4. What predicts participants' choices?

Following the aforementioned analyses, we attempted to predict individual choices in the mouse-tracking task based on their scores at the side tasks. We built a Poisson generalised linear model (since choices were count data) entering Complexity Score, Unpredictability Score, Inhibition Score and Shifting Score as fixed effects.

In the full sample, the full model could not significantly predict External Reward choices ($F(1,8) = 2.074$, $p = .979$) nor Informational Reward choices ($F(1,8) = 1.789$, $p = .987$), but it did predict Novel Stimulation choices ($F(1,8) = 21.808$, $p = .005$). Specifically, Working Memory Score was a significant predictor ($p = .002$): for every extra level completed in the WM task, 1.29 fewer NS choices were made (95% CI, -.213 to -.046).

Within each age group, the model did not significantly predict choices. Specifically, in the children groups the full model could not significantly predict External Reward choices ($F(1,7) = 1.644$, $p = .977$), Informational Reward choices ($F(1,7) = 5.292$, $p = .624$), nor Novel Stimulation choices ($F(1,7) = 7.450$, $p = .384$).

In the adolescents group, the full model could not predict External Reward choices ($F(1,8) = 2.407$, $p = .966$), Informational Reward choices ($F(1,8) = 3.030$, $p = .932$), nor Novel Stimulation choices ($F(1,8) = 7.258$, $p = .509$).

Finally, in the adults group, the model did not predict External Reward choices ($F(1,7) = 6.493$, $p = .484$), however Uncertainty preference seemed to be a marginally significant predictor ($p = .051$). Specifically, it seemed that participants who preferred the intermediate uncertainty option in the relevant task chose the ER option significantly less (-.600 ER choices, 95% CI, -1.092 to -.109, $p=.017$). Otherwise, the model did not predict the Informational Reward choices ($F(1,7) = 2.958$, $p = .889$), nor Novel Stimulation choices ($F(1,7) = 8.072$, $p = .326$).

3.3.2.5. Summary of Results

We observed significant differences in participants' choices depending on their age. Overall, children did not differ in their choices throughout the task, making balanced decisions between the three options. Compared to the other groups, they chose the NS option significantly more than adolescents and adults. Adolescents and adults, on the other hand, did not differ in terms of their ER and IR choices, and both chose ER and IR more often than the NS option.

When participants' choices up to the ER, IR or NS completion were analyzed, we found that children chose the IR less than adults and the NS more than adults, whereas the groups did not differ in other choices. Children also did not show any difference between their choices, while adolescents chose ER more than the other options. Adults chose ER and IR more than the NS option.

Regarding movement parameters, the type of dilemma significantly affected trajectory types. Specifically, in all age groups, participants' trajectory paths in the IR-NS dilemma showed a discrete change of mind more often than in the ER-NS dilemma. The type of dilemma did not affect trajectory curvature measures (MAD, MDabove and AUC), nor RTs and x-flips. Age groups significantly differed in their movement complexity. Specifically, children made more x-flips, and their trajectories followed a straight path significantly more often than adults.

In the complexity preference task, participants in all age groups chose the less complex option significantly less than the other three, which did not differ in terms of frequency. In the unpredictability preference task, only adults chose the more unpredictable option more than the less unpredictable, while the other groups did not show any significant differences in their choices. In the Go/No-Go and Memory tasks, children performed significantly worse than adolescents and adults, who reached ceiling performance. In the Switching task, no significant differences were observed between groups.

Lastly, Working Memory scores significantly predicted the NS choices in the regression model when all age groups were analysed together. Specifically, lower Working Memory scores predicted more NS choices. No other score successfully predicted choices in the mouse-tracking task. Within the age groups, preference for Intermediate Uncertainty in the adults group marginally resulted in fewer ER choices.

3.3.3. Discussion of Experiment 2

Our second experiment replicated some of the findings of the first one, but also differed in important aspects. Regarding participants' choices, all age groups were interested in acquiring the external reward, but offering them more tokens (trials) to spare made them choose the other options for a significant amount of times. Specifically, adults followed the same (optimal) strategy with Experiment 1, acquiring the external reward and also achieving the (implicit) learning goal. Adolescents followed the same strategy, since the design of this experiment forced them to keep making choices even after the acquisition of the external reward. Indeed, when the choices they made up to the ER acquisition were analysed, it became obvious that they followed the same ER-driven strategy as in Experiment 1. In contrast, children preferred to balance their extra choices between the two exploratory options, choosing the novel stimulation option an equal number of times with the informational reward option. This finding is consistent with some previous findings, which show that school-aged children still do a lot of random exploration (Meder et al., 2019), as well with findings showing that longer decision horizons favour exploratory choices in general (both directed and random). However, other findings suggest that children in this age already choose a lot based on information gaps and in a (systematic) way that supports their learning progress (Blanco & Sloutsky, 2021; Schulz et al., 2019). This employment of different strategies in children might be explained by individual differences in a large spectrum, as well as within-subject differences per block or even trial (e.g., Siegler, 2007, discusses the change of strategies as children learn to do a task or change goals trial-to-trial). Inspecting the individual strategies in both of our experiments showed that some children preferred to choose based on IR, while others consistently preferred NS, at least for some of the blocks, suggesting that cognitive or dispositional factors might play a role in these differences. Thus, we measured participants' executive functions and preferences towards complexity and unpredictability to shed some light to these factors.

Our results suggest that some of these factors might play a significant role in participants' individual exploratory strategies, although these findings should be interpreted carefully. We had hypothesised that cognitive control (inhibition being the more possible candidate) and preference for unpredictability will play an important role in participants preference for the novel stimulation choices (Gopnik, 2020; Litman, 2010). From our findings, only WM score was a significant (negative) predictor of NS choices. However, as this was probably the task which captured age differences in the more accurate way, this finding might just reflect the (already observed) differences in NS choices between the groups, and not have any explanatory value. Furthermore, certain design decisions that we made might have influenced how accurately we measured the aforementioned variables. Specifically, although we captured age-related differences in the Go/No-Go task, it is possible that we did not capture individual differences in inhibitory ability (which would be the important ones in our regression analysis, as we were interested in predicting individual scores). The reason might have been the decrease in experimental trials per block (compared to the Bezdjian et al. (2009) study), which might have not been sufficient to establish a strong prepotent Go response, and would thus make the task much easier for more proficient participants (especially for adolescents and adults). Furthermore, our preference for unpredictability task also included very few trials (20 per block) and might thus have made the statistics of each option hard to learn – and as a result, prefer – for younger participants.

The variables of our mouse-tracking task might also have been confounded. Specifically, adolescents' and adults' choices show a clear preference for IR options compared to NS ones, which can be motivated by closing information gaps and obtaining information based on previous knowledge and informational uncertainty (a behaviour, as stated at the introduction, close to directed exploration). However, the fact that a puzzle implicitly has a specific goal (its completion), it is unclear whether participants were drawn by trial-by-trial informational value, or by the overall aim to achieve a goal – and similarly experience satisfaction by it. Even on a trial-by-trial basis, taking rewarding steps to achieve a long-term goal has been shown to be intrinsically motivating (Woolley & Fishbach, 2017), in a different way compared to knowledge. Furthermore, although our NS option involved the presentation of new cartoon (or animal) images on every trial, the category of the images was predictable on every block, e.g., images of horses would appear. This might have decreased the amount of novelty, while at the same time increasing the possible effect of individual interests (e.g., someone might dislike horses, and thus choose more the IR option on that block). As a result, we decided to address these issues by making changes in the task in our next experiment.

Finally, although we had aimed to capture the decision process more in the movement parameters by imposing a time-limit in our mouse-tracking task, our results showed no differences in most trajectory features between dilemma types, in contrary to Experiment 1. This finding might be explained by this time-limit, as recent research suggests that time-limit might specifically influence the initial stages of the movement, and thus affect the curvature features more (e.g., Wirth et al., 2020, suggest that free initiation should be preferred if curvature measures are of interest). However, our finding might be explained by the differences in the task itself, and specifically by the increase of the number of trials, which might have added less pressure and conflict on individual decisions (since many options could be chosen across a block).

Thus, to assess these matters, we proceeded in our third experiment.

3.4. Experiment 3

In Experiment 3, we decided to examine children only, as we were mainly interested in refining our experimental manipulations of informational reward and novelty and replicate our previous findings with a slightly modified task. We specifically tested 5- to 7-year-olds (as in Experiment 2) and overall kept the experimental design identical, implementing small changes in the mouse-tracking task procedure (i.e., removing the time limits) and, most importantly, in the task stimuli. Regarding the decision horizon, we decided to keep a long horizon (i.e., participants could still make choices after they had acquired the ER), so that we could compare our findings with those of Experiment 2. Lastly, we had participants complete three EF tasks (inhibition, working memory and switching abilities), different from those used in Experiment 2, in order to capture individual differences more accurately.

We expected to replicate our previous findings; i.e., that children will be equally driven by the three options and balance their choices from the beginning of the block. We also expected mouse-tracking measures to reflect the conflicts more accurately, since we returned to a time-limitless procedure as in Experiment 1. Eventually, we expected the EF performance to correlate with ER and NS choices, such that better EF performance will favour exploitatory choices.

3.4.1. Methods

3.4.1.1. Participants

In this experiment, the participants were 18 children (9 females, mean age: 6.09 years). All participants were neurotypical and had normal or corrected-to-normal vision. The participants were volunteers recruited through the Birkbeck Babylab database and local primary schools. They received a gift for their participation.

3.4.1.2. Mouse-tracking task

3.4.1.2.1. Design

Participants had to complete four experimental blocks of 27 trials each. In each trial, two of the total three options (selected randomly) appeared on the screen. The participants had to choose an option (in a *2 alternative forced choice* format) in one of the following dilemmas: External Reward VS Informational Reward (ER-IR), External Reward VS Novel Stimulation (ER-NS), or Informational Reward VS Novel Stimulation (IR-NS). If the External Reward (ER) option is chosen for a total of nine times, the participant won the reward. The trials continued until a total of 27 trials was completed. This means the maximum choices for each option could be 18.

3.4.1.2.2. Stimuli and procedure

The general layout of the decision-making task remained the same as in the previous two experiments: the two alternative options appeared in fixed positions on the top of the screen during the test trials (or in random positions during the training trials) and participants had to click and drag a third icon (a heart) on top of the options make a choice. However, the stimuli themselves differed (Figure 3.19). Specifically, the three dogs were removed and the decision screen could now show the following options:

- For the ER, it depicted the tower on its current stage of completion, e.g., if the ER option was already chosen in four previous trials, the tower appeared with four floors already built. If this option was chosen, the next screen (reward screen) depicted the tower with an extra floor added on top.

- For the IR, it depicted a blurry picture. The pictures were chosen from the large dataset used by Jepma, Verdonschot, Van Steenbergen, Rombouts, & Nieuwenhuis (2012; shared with us by the authors). They had a resolution of 71 dpi, and were centered on a white rectangle of 197 × 281 pixels. The blurred version of each picture was created by Jepma et al. (2012) by means of Gaussian smoothing with a radius of 20–22 pixels. When this IR option was chosen, the reward screen revealed the clear version of the picture. No picture was shown more than once.
- For the NS, the option was a red question mark, indicating that an unknown image would follow. Indeed, if it was chosen, an image was shown in the reward screen, which was drawn randomly from a large variety of images : they varied in terms of type (e.g., realistic photo, drawing, sketch, cartoon) and theme (e.g., animals, landscape, fruit, people, household items). These images were chosen on the basis of being highly unpredictable, while the images in Experiments 1 and 2 were still new on each trial but belonged to the same category (animals).

Procedure was almost identical to Experiment 2, with the following important change applied both on the training and the test phase: participants did not have a time limit to start the trial (i.e., to move the heart) or to make the choice after they had started their movement. However, as in Experiment 2, the two options did not appear unless the bone had been moved a minimum of 4 pixels above its initial position.

Figure 3.19

The new dilemmas and stimuli. Above, the dilemmas, left to right: ER-NS, IR-NS, ER-IR. Below, the reward screens, left to right: NS, IR, ER.

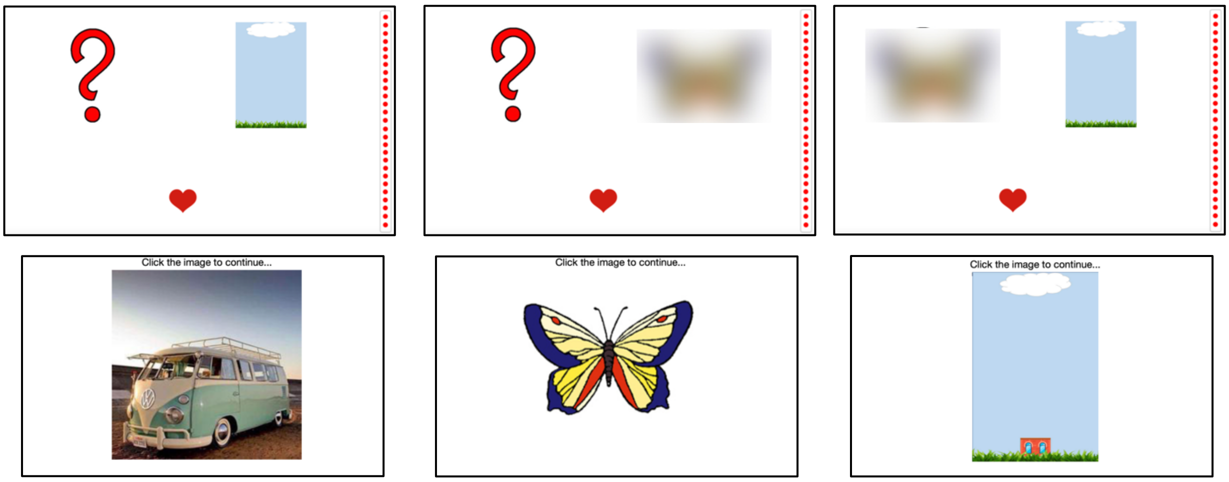

3.4.1.3. Executive functions tasks

3.4.1.3.1. Switching task

We used a child-friendly version of a switching task previously developed in our lab (Carteron, 2022), where both the sets of stimuli and responses changed simultaneously, with the classification dimension being intrinsically determined by the stimuli. A similar version of this task designed for adults (Rogers & Monsell, 1995) asks participants to classify numbers as either odd or even (task set 1) and letters as vowels or consonants (task set 2). In our case, the task sets involved categorizing animals as either sea or land creatures and categorizing objects as sports or food items. Consequently, the classification dimension was dictated by the nature of the stimulus, be it an animal or an object. The task was administered using a laptop and children indicated their choices by pressing either the left or right key on a keyboard. There was no time limit in their response. The mapping of left and right keys was represented using pictograms of sea and land (for task set 1) or sports and food (for task set 2), displayed on the left and right sides of the screen, as illustrated in Figure 3.20. This setup ensured that when the central icon represented an animal, a sea pictogram consistently appeared on the right, with a land pictogram consistently on the left. In contrast, when the central icon depicted an object, a sports pictogram was consistently displayed on the right and a food pictogram was consistently shown on the left. A total of six distinct animals and six distinct objects were used. At the beginning of the task, practice trials were completed with all of the 6 possible stimuli, and ensured instructions were understood. Then a block of 10 trials of the animal task set was presented, followed by 10 trials of the objects task set (Block 1 and 2; non-switch trials), followed by a block of 20 trials alternating between each task set (Block 3; switch trials). We measured response times in switch and non-switch trials, as it is considered a more informative measure compared to accuracy, representing the extra processing cost in switch trials. The final score was calculated as the difference between the two scores.

Figure 3.20

Example of a display of the switching task. Participants had to categorise the animal as a "land animal" or a "sea animal".

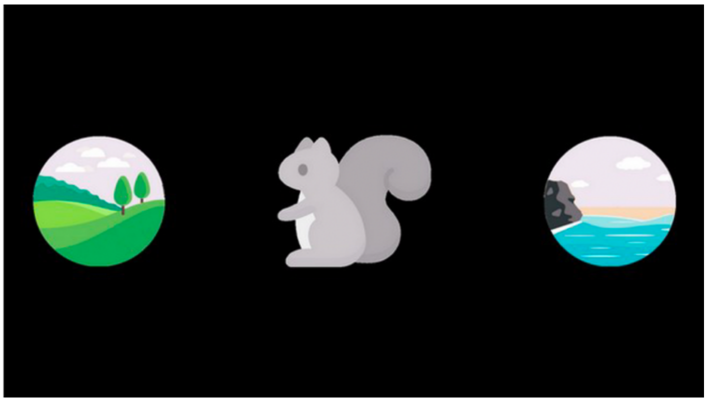

3.4.1.3.2. Inhibition task

Inhibition ability was measured using the BAT task (Figure 3.21), a child-friendly version of the go/no-go task designed by Schröer, Cooper and Mareschal (2021) and based on similar tasks (e.g., Drechsler, Rizzo, & Steinhausen, 2010; Sobeh and Spijkers, 2013). The task was administered using a laptop. Children were asked to respond by pressing the space bar every time they saw a bat appear on the screen, because bats can turn into vampires, but instead refrain from pressing the bar when they saw a cat, because cats are good. They were instructed to press the space bar as fast as possible when they saw the bat ("go trials"), but not when they saw a cat ("no-go trials"). In the beginning of the task, children completed two practice trials, one for each animal, and they had to answer what they were supposed to do when they saw each animal. To ensure that the go response was the default action, the children were initially presented with five consecutive go trials. The remaining trials were presented in a random order. The majority of the trials, specifically 74% of the total 35 trials, were go-trials. Children were given a 2-second window in which to respond before the image disappeared from the screen, and the subsequent trial commenced following a 1-second interstimulus interval. The inhibition score was computed by dividing the errors (false alarms/commission errors), by the total number of trials.

Figure 3.21

Example of the trial sequence in the BAT task. Participants had to press Space when they saw a bat, but not when they saw a cat.

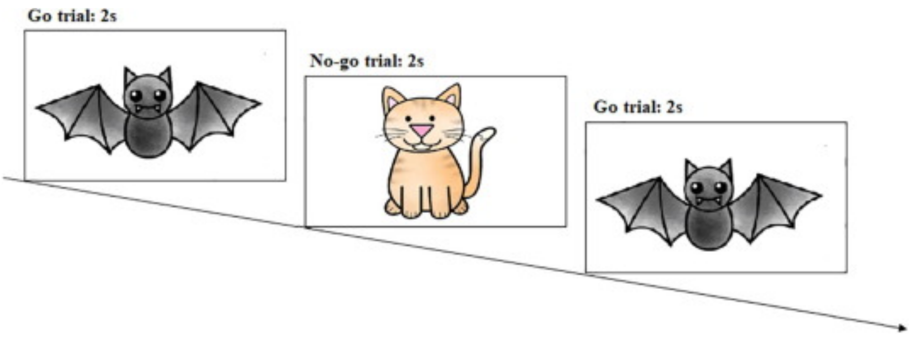

3.4.1.3.3. Working Memory task

Working memory was measured using the original version of the "Not this" task from the Early Years Toolbox (EYT; Howard, & Melhuish, 2017), which was completed on a tablet. We also used the EYT 'Not This' task (which measures phonological WM). In this task, as in the one used in Experiment 2, children had to follow auditory instructions of increasing complexity (Figure 3.22). Children are asked to choose a stimulus that is *not* of a particular colour, shape or size (or some combination of these). The task includes five trials at each level of complexity (levels 1 to 8), whose difficulty depends on the number of stimulus characteristics that must be simultaneously activated in mind (i.e., level 1 involves maintaining one feature in memory, level 2 involves maintaining 2 features in memory, and so on). Each trial has the following procedure: (1) an auditory instruction is played against a white screen; (2) a 3 second delay follows against a white screen; and then (3) a 4 x 5 array of different shapes of different colours and sizes with cartoon faces are presented until a response is made by tapping the shape(s) that the participant believes match to the instruction. The task continues until the full completion (at level 8, eight features to remember) or failure to accurately complete at least three of the five trials within a level. Performance is measured using a point score, which is computed as: starting from level 1, one point for each level in which at least three of the five trials were performed correctly, plus 1/5 of a point for all correct trials thereafter.

Figure 3.22

Array of target shapes in Working Memory task. Participants had to remember auditory instructions about colour/shape/size combinations to pick to correct shape(s).

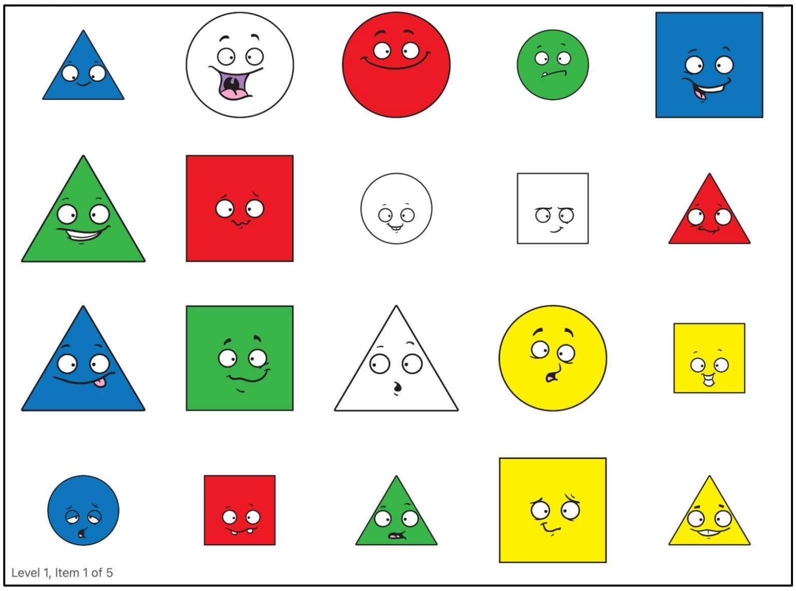

3.4.2. Results

One participant was excluded from the study as they needed help to with the mouse to complete the mouse-tracking study, and thus the movement data could not be used. Outliers in the separate tasks were excluded from the relevant analyses. The exclusion criteria will be discussed separately below for each task.

3.4.2.1. Choices

Figure 3.23a and Table 3.8 show the mean number of choices and the SDs of each of the three options. We initially analysed the choices that participants made and found a significant main effect of Option: $F(2, 34) = 5.682$, $p = .007$. Pairwise comparisons showed that in total, children chose the ER option significantly more than the NS option ($p=.011$) whereas the ER and the IR option did not differ ($p = .092$), neither did the IR and the NS option ($p = .068$).

Table 3.8

Means and SDs of choices per Option

	Mean choices (sd)
External Reward	9.592 (2.106)
Informational Reward	8.197 (3.079)
Novel Stimulation	6.671 (2.926)

Figure 3.2*33*

Number of choices per Option in the Curiosity task, a) until block completion, b) until first option completion

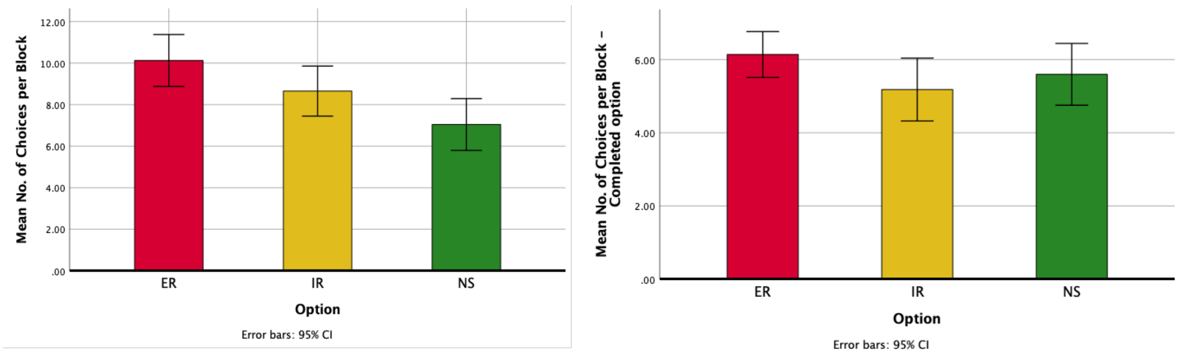

a) b)

To examine if participants were affected by the longer decision horizon, we also analyzed their choices up to the completion of at least one of the options (Figure 3.23b). We found no significant effect of Option $F(2,34) = 1.322$, $p = .280$.

3.4.2.2. Mouse-tracking data analyses

The same procedure as in Experiments 1 and 2 was followed to extract features from the raw mouse-tracking data. Table 3.9 shows means and SDs for the five indices we chose to compare for each dilemma.

Table 3.9

Means and SDs for movement parameters per Dilemma and Age Group

		Children	
Features[a]		Mean	SD
Maximum Absolute Deviation (MAD)	ER-IR	127.99	151.84
	ER-NS	112.19	138.67
	IR-NS	126.77	152.47
Maximum Deviation above ideal line (MDabove)	ER-IR	141.32	135.571
	ER-NS	125.06	123.92
	IR-NS	140.77	135.98
Area Under the Curve (AUC)	ER-IR	29542.90	57440.08
	ER-NS	28639.84	43080.93
	IR-NS	27859.48	47916.56
x-flips	ER-IR	1.72	1.61
	ER-NS	1.60	1.43
	IR-NS	1.87	1.69
Response Times	ER-IR	3980.31	1965.11
	ER-NS	3818.76	1665.34
	IR-NS	4062.83	1977.19

[a]All time related values are presented in milliseconds (ms), all position related values are presented in pixels (px), area (AUC) is displayed in px^2.

Our analyses showed that the type of Dilemma did not affect MAD ($\chi^2(2) = 3.585$, $p = .166$) or MDabove ($\chi^2(2) = 4.698$, $p = .095$. Furthermore, there was no significant effect of type of Dilemma on AUC ($\chi^2(2) = 0.375$, $p = .829$). There was a significant main effect of Dilemma on x-flips ($\chi^2(2) = 6.495$, $p = .039$; Figure 3.24). However, pairwise comparisons showed no significant differences between the ER-IR and ER-NS dilemmas ($p=.426$) or between the ER-IR and IR-NS dilemmas ($p=.583$). A marginally significant difference was found between the ER-NS and IR-NS dilemmas ($p=.054$), specifically showing that children made more flips in the IR-NS dilemma. Lastly, no significant main effect of Dilemma was found for RTs ($\chi^2(2) = 4.971$, $p = .083$).

Figure 3.24

Mean number of x-flips per type of Dilemma.

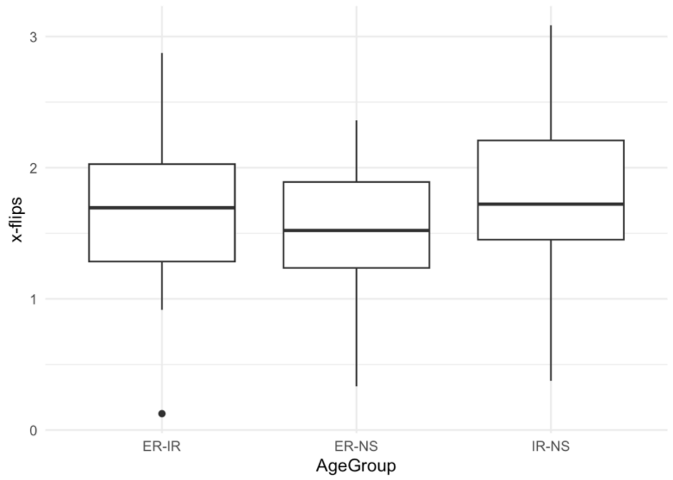

Figure 3.25 shows the automatically clustered trajectories. There was a significant effect of Dilemma on trajectory types ($\chi^2(2) = 16.413$, $p < .001$), and specifically the IR-NS dilemma was significantly different from ER-NS ($p<.001$) but not from ER-IR ($p = .100$). The ER-IR and ER-NS dilemmas did not differ ($p=.115$). To further investigate these differences, we compared each type of trajectory per type of Dilemma separately (Table 3.10). We found that the participants significantly differed in the frequency of straight trajectory paths in each dilemma ($F(1.513, 25.725) = 8.819$, $p = .003$), and more specifically, they had significantly more straight paths in the IR-NS dilemma when compared to ER-IR ($p=.035$) or to ER-NS ($p = .011$). Also, participants significantly differed in the frequency of curved trajectory

paths in each dilemma (F(2,34) = 4.150, p =.024), but pairwise comparisons did not show any significant difference – IR-NS dilemma had marginally less curved trajectories than the ER-NS one (p=.058). Finally, there was significant difference in their trajectories which showed a continuous change of mind (F(2,34) = 3.805, p = .032, but pairwise comparisons did not show any significant difference.

***Table 3.**11*

Frequencies of trajectory types per Dilemma.

	Straight	**Curved**	**cCoM**
ER-IR	515	70	13
ER-NS	490	84	20
IR-NS	546	51	10

Figure 3.25

All observed trajectories as they have been automatically clustered in three different prototypes by mousetrap. No dCoM and dCoM2 types were observed.

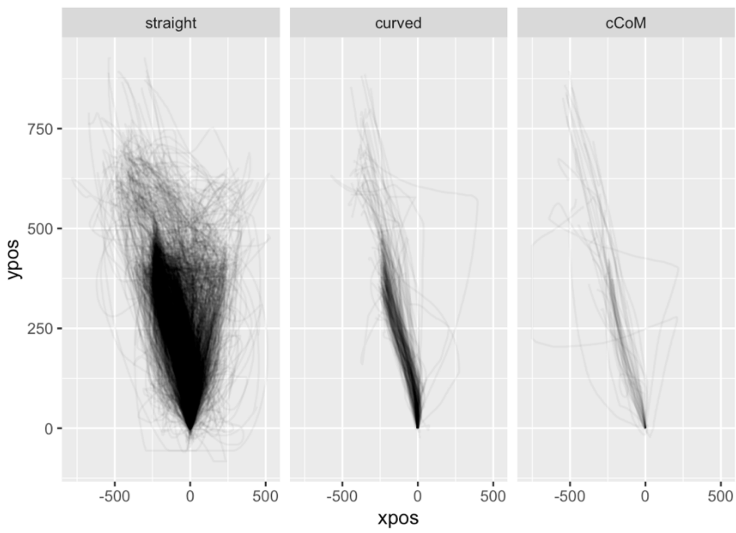

3.4.2.3. Secondary tasks

Participants' performance (means and SDs) in the EF tasks is shown in Table 3.11. Go/No-Go scores are participants' commission-error rate (i.e., inability to inhibit response in No-Go trials), Switching score is participants' RT difference in mixed vs. simple blocks (i.e., the extra processing cost in blocks where rule switching was required) and working memory scores reflected completed levels and trials within the task.

Table 3.12

Participants' performance in EF tasks

Task (measure)	*Mean score*	*SD*
Inhibition (error rate)	0.160	0.133
Switching (ms)	102	31
Working Memory (level)	2.933	0.172

3.4.2.4. What predicts participants' choices?

We also tried to predict individual choices in the mouse-tracking task based on their scores at the side tasks. We built a Poisson generalized linear model entering Inhibition Score, Switching Score and Memory Score as fixed effects.

The full model could not significantly predict External Reward choices ($F(1,3) = 2.535$, $p = .469$), Informational Reward choices ($F(1,3) = .948$, $p = .814$), nor Novel Stimulation choices ($F(1,3) = 2.018$, $p = .569$).

3.4.2.5. Summary of Results

We observed significant differences in participants' choices. Overall, children did not differ in terms of their ER and IR choices, or in terms of their IR and NS choices, but they chose ER more often than the NS option. When only their choices up to the initial ER, IR or NS completion were analysed, we found that children did not show any difference between their choices, balancing the three options.

Regarding movement parameters, the type of dilemma marginally affected trajectory types. Specifically, children's trajectory paths in the IR-NS dilemma were straight more often than in the other dilemmas, and curved less often than in the ER-NS. The type of dilemma did not affect trajectory curvature measures (MAD, MDabove and AUC), nor RTs. Children made significantly more x-flips in the IR-NS dilemma compared to the ER-NS dilemma, however this difference was only marginally significant.

Finally, none of the individual EF scores predict choices made in the regression model.

3.4.3. Discussion of Experiment 3

The aim of Experiment 3 was to replicate our previous findings, while measuring the different variables (ER, IR and NS) in a more accurate way by unconfounding them from possible factors that might have influenced decisions in the previous experiments. More specifically, we were interested in capturing behaviours driven by IR as a "trial-by-trial" relief of perceptual curiosity, rather than a longer term information gathering goal spread over several trials, and NS in a set of stimuli which could not be predicted – and thus being consistently novel.

Our results replicate some of our previous findings. Children seem to prefer the ER choices compared to the NS ones as a general tendency, a finding that replicates the findings of our first experiment but contradicts the findings of the second. One possible reason for this difference could be the time limit imposed on decisions in Experiment 2 that might have forced participants to choose more impulsively – and thus maybe taking more into account the immediately rewarding option compared to a delayed rewarding state. Of course, we cannot rule out the possible effect of the new NS stimuli, which were less predictable – and this amount of novelty might have created a certain amount of aversion to children, or at least some of them – stable or contextual predispositions towards unpredictability might play such a role, as was previously discussed in Experiment 2.

Furthermore, the new EF tasks performance did not predict participants' choices, a finding consistent with Experiment 2. This was also the case for WM. This suggests that individual differences in choice selection are not explained by cognitive factors as we had initially hypothesised, but might instead be related to attitudes and traits, such as trait curiosity and tolerance for uncertainty.

The movement trajectories revealed some differences in the IR-NS dilemma. This was also the case in our previous experiments. Interestingly, in this experiment we found straighter trajectories, and more x-flips in the IR-NS conditions as compared to the other dilemmas, a movement profile which seems contradictory. However, if we take into account the categorisation of trajectory types, to continuous vs. more discrete paths, it seems that this dilemma probably did not have a continuous conflict like the ER-NS dilemma (since it had significantly fewer curved paths), but more likely had small micro-changes of direction, possibly too short to be categorised as discrete changes of mind, but still flips on the y-axis. The interpretation of this difference is not straightforward, as a similar conflict could have been observed in the ER-IR dilemma if it was just a reflection of ongoing conflict. nevertheless, as in our previous experiments the IR-NS dilemma seemed slightly more difficult for children.

The lack of other differences in the curvature or RT measures, even without a time limit, could be explained by the temporal horizon of decisions, as in Experiment 2. Because participants knew that they had enough trials to both explore and obtain a reward, the pressure on each decision became smaller.

Finally, an important point should be made regarding the value comparison between the options. Even though we replaced the IR goal option with trial-by-trial perceptual uncertainty resolution, the ER option still consisted of a goal – participants had to choose this option repeatedly in order to finally acquire the reward, but they were not explicitly rewarded on each trial – unless only by association. In an attempt to make the values of the options even more comparable in terms of instant reward, in our next experiment we decided to change the ER option content.

3.5. Experiment 4[4]

In our final experiment, we were interested in dissociating the value of the external reward from the goal-following process. It is known that decision-making is often influenced by the goals that people choose to pursue, which might take into account the subjective value of available options, but also influence the valuation of these options in return (Frömer & Shenhav, 2022; Molinaro & Collins, 2023). For example, in our paradigm, higher order preferences such as deciding to complete the

[4] Data collection for experiment 4 was carried out by Dian Qu, graduate student in Birkbeck MSc Psychology, as part of her final dissertation.

ER goal for the sake of it (or because they are told to) might explain participants' behaviour in a specific block, and as a result their preference for this option might not reflect how rewarding they actually find it. In this case, we would not be comparing between reward values, but separate goals. As a result, in this experiment we replaced the ER option with stimuli associated with primary or secondary rewards (food or money) which could be gained on every trial, but had no end goal. No other changes were made to the procedure described in Experiment 3. We also decided to only focus on two executive function skills in this experiment (inhibition and switching), as we considered them better possible candidates for an association with exploratory tendencies. Our participants were 5- to 7-year-olds, 13- to 15-year-olds and adults.

We expected participants to have similar preferences for the ER option as in previous experiments; i.e., to still choose ER either equally to the other options (children) or significantly more than NS (adolescents and adults).

3.5.1. Methods

3.5.1.1. Participants

In this experiment, the participants were 20 children (9 females, mean age: 6.32 years), 10 adolescents and 20 adults. All participants were neurotypical and had normal or corrected-to-normal vision. The participants were volunteers recruited through the Birkbeck Babylab database, word of mouth and local primary schools. They received a gift for their participation.

3.5.1.2. Mouse-tracking task

3.5.1.2.1. Design

Participants had to complete four experimental blocks of 27 trials each. In each trial, two of the total three options (selected randomly) appeared on the screen. The participants had to choose an option (in a *2 alternative forced choice* format) in one of the following dilemmas: External Reward VS Informational Reward (ER-IR), External Reward VS Novel Stimulation (ER-NS), or Informational Reward VS Novel Stimulation (IR-NS). The trials continued until a total of 27 trials was completed. This means the maximum choices for each option could be 18.

3.5.1.2.2. Stimuli and procedure

The general layout of the decision-making task remained the same as in the previous two experiments: the two alternative options appeared in fixed positions on the top of the screen during the test trials (or in random positions during the training trials) and participants had to click and drag a third icon (a heart) on top of the options make a choice. However, the stimuli for the ER option differed (Figure 3.26a-d). Specifically, instead of the tower image, on the decision screen children were now seeing a bag of candies. If they chose this option, the reward screen depicted three candies and a stack of candies that was getting one candy taller every time this option was chosen. Instead, adolescents and adults saw a bag with money on the decision screen. If they chose this option, the reward screen showed them three coins and a stack of coins getting on coin taller each time. These stimuli were chosen as candies and money are commonly used (external) rewards in armed-bandit tasks – and gambling games in general.

The rest of the procedure was identical to that of Experiment 3.

Figure 3.26

Example decision and reward screens for children (a,b) and adolescents-adults (c,d)

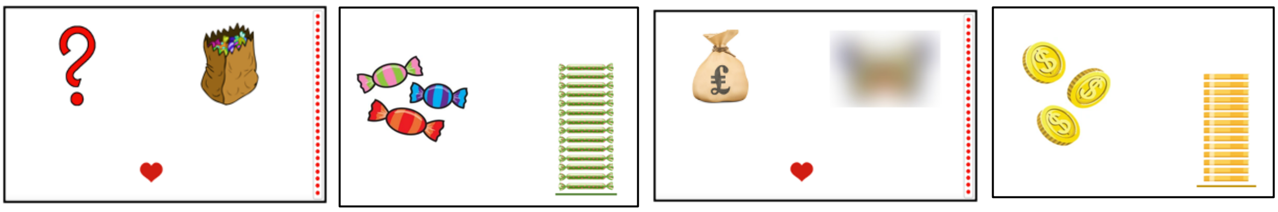

a) b) c) d)

3.5.1.3. Executive functions tasks

3.5.1.3.1. Inhibition task

For children, we used the BAT task as in Experiment 3.

For adolescents and adults, we used a simple Go/No Go task, similar to the one used in Experiment 3. Participants were instructed to press the Space key on their keyboards when the letter P appeared among three stars, in one of four random positions in a square grid (Go Stimulus; Figure 3.27), but they had to withhold their response when the letter R appeared instead (No-Go stimulus). Each trial started with a presentation of a block of four stars for 1500ms (Figure 8c), followed by the target stimuli for 500ms. The response screen showed four stars again for 1200ms. Participants had to respond during this time limit, or the trial was recorded as an error. Participants did not receive any feedback for their responses, and the next trial commenced immediately after the response/time limit. The block started with 5 go-trials to create a prepotent Go response and the majority of the trials (30 out of 40 trials) were go-trials. The inhibition score was the false alarm rate, which was computed as the ratio of NoGo trials in which a response was given by the ones it was inhibited (commission errors).

Figure 3.27

Example grid of Go and No-Go trials in adolescents' and adults' Inhibition task

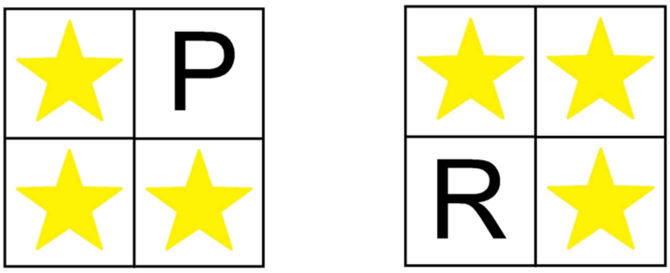

3.5.1.3.2. Switching task

For children, we used the same task as in Experiment 3.

For adolescents and adults, we used a standard Dimensional Change Card Sort task (Zelazo, 2006). Participants were asked to sort coloured shapes (red and blue triangles and squares) according to one of the two dimensions: colour or shape (Figure 3.28). On each trial, participants would always see a red triangle and a blue square on the bottom of their screen, always with the same right-left mapping. A word would appear between the shapes for 500ms, indicating the sorting rule: COLOUR or SHAPE), and the target shape would then appear in the middle of the screen until a response was given. Participants were asked to press "A" when the target shape was red or a triangle, and "L" when the target shape was blue or a square. An interstimulus interval of 800ms followed the response. There were 20 shape and 20 colour trials, and 20 mixed trials. We measured response times in switch and non-switch trials, and the final score was calculated as the difference between the two scores.

Figure 3.28

Example trials for each sorting rule. Participants had to categorize the middle shape according to the shown rule. The trials could appear in separate or in the same (mixed) blocks.

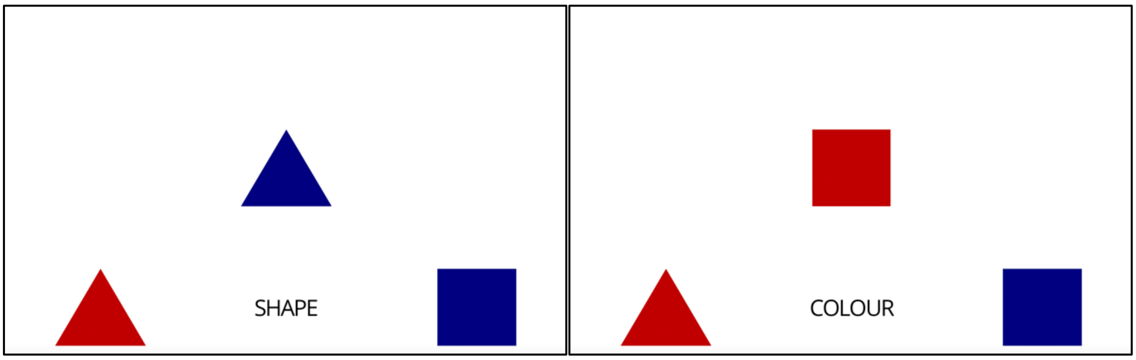

3.5.2. Results

No participants were excluded from this study. Outliers in the separate tasks were excluded from the relevant analyses. The exclusion criteria will be discussed separately below for each task.

3.5.2.1. Choices

Figure 3.29 and Table 3.12 show the mean number of choices and SDs of each of the three options for each age group. We first analysed the choices each age group made and found no significant interaction between Age Group and Option: $F(3.532, 83.004) = 1.180$, $p = .325$, Greenhouse-Geisser corrected. There was no main effect of Option ($F(1.766, 83.004) = 1.392$, $p = .254$, nor Age ($F(2,47) = 2.883$, $p = .066$. Since it was part of our hypotheses, we also inspected the pairwise differences between age groups' choices. Adults chose the NS option significantly more than children ($p = .029$), but the three groups did not differ significantly in the other choices. We did not observe any significant differences in the choices within each age group either.

Table 3.13

Means and SDs of choices per Option for each age group

	Children	*Adolescents*	*Adults*
External Reward	10.150 (5.850)	10.759 (4.760)	8.737 (6.575)
Informational Reward	8.100 (3.159)	8.575 (3.790)	9.713 (5.441)
Novel Stimulation	6.050 (4.080)	7.675 (3.817)	9.900 (5.173)

Figure 3.29

Number of choices per Option for each age group in Curiosity task

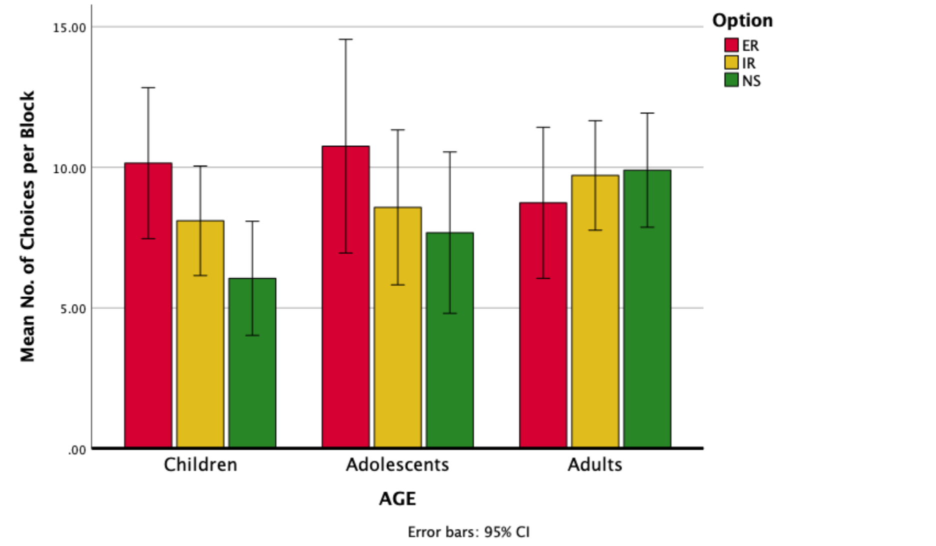

3.5.2.2. Mouse-tracking data analyses

We followed the same procedure as in the previous experiments to extract features from the raw mouse-tracking data. Table 3.13 shows means and SDs for the five features we chose to compare between each dilemma. The unfolding of aggregated trajectories across time steps for each Dilemma can be seen at Figure 3.30.

Figure 3.30

Change of mouse position on x-axis during the progress of each trial (time-normalised trajectories). Data have been aggregated for each dilemma. (a). Different unfolding of decisions over time can be observed in children (b), adolescents (c) and adults (d)

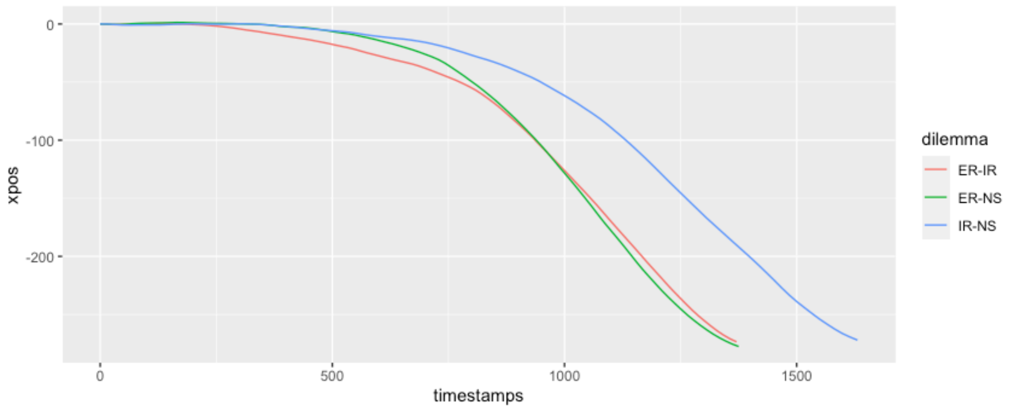

a)

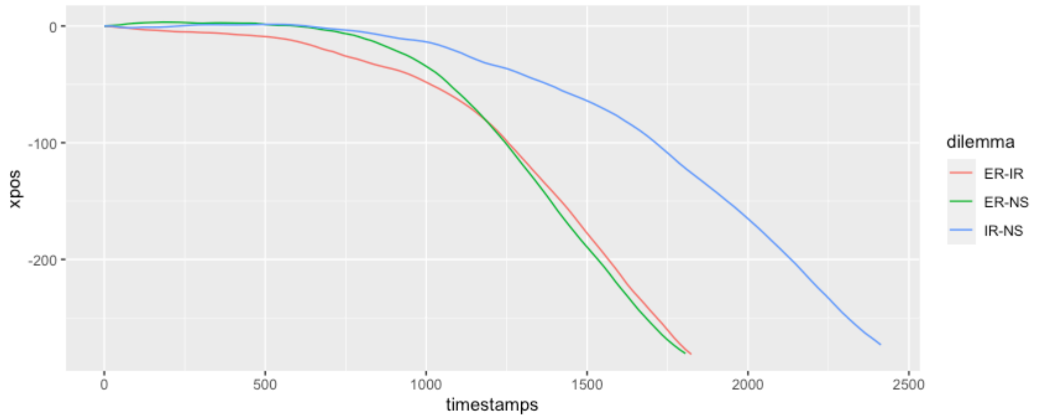

b)

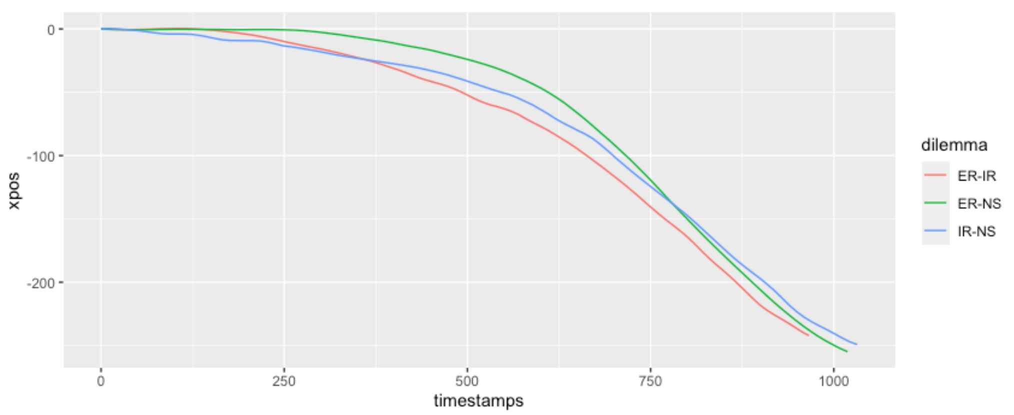

c)

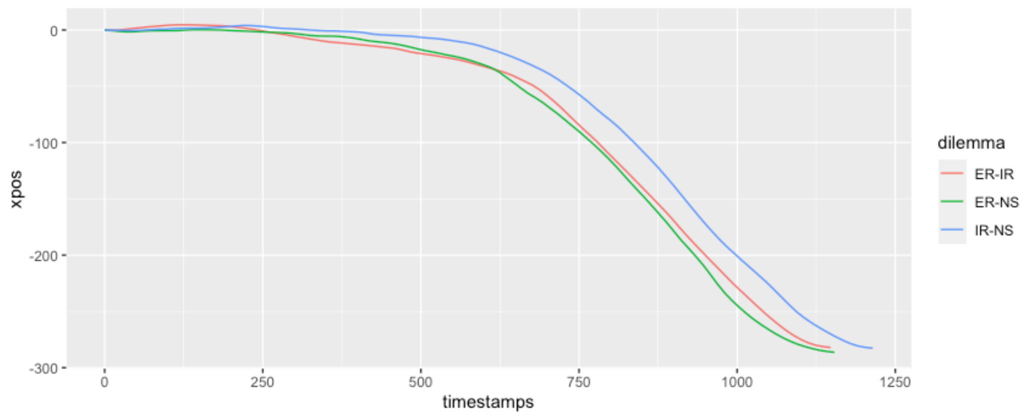

d)

Table 3.14

Means and SDs for movement parameters per Dilemma and Age Group

		Children		Adolescents		Adults	
Features[a]		Mean	SD	Mean	SD	Mean	SD
Maximum Absolute Deviation (MAD)	ER-IR	107.04	137.49	84.37	130.63	124.41	138.69
	ER-NS	118.01	143.64	113.43	141.77	116.08	140.82
	IR-NS	133.40	155.71	91.61	135.42	122.05	138.82
Maximum Deviation above ideal line (MDabove)	ER-IR	122.34	116.37	99.97	111.69	130.09	131.64
	ER-NS	131.99	123.34	123.79	130.03	123.44	132.99
	IR-NS	150.95	133.57	104.23	121.96	127.60	132.50
Area Under the Curve (AUC)	ER-IR	26668.42	34503.86	23932.44	40367.07	34836.22	44462.16
	ER-NS	29838.19	41123.82	29443.65	39185.82	31194.98	41654.78
	IR-NS	29033.79	37930.33	26227.24	108296.89	34818.59	46180.97
x-flips	ER-IR	1.91	1.86	1.17	1.13	1.42	1.34
	ER-NS	2.06	2.29	1.31	1.41	1.46	1.39
	IR-NS	2.71	2.57	1.31	1.49	1.40	1.27
Response Times	ER-IR	3737.29	2434.34	2257.51	941.61	2740.68	1294.38
	ER-NS	3798.59	3160.80	2272.82	880.14	2696.49	1327.85
	IR-NS	4298.20	3301.51	2308.67	925.59	2796.44	1318.56

[a]All time related values are presented in milliseconds (ms), all position related values are presented in pixels (px), area (AUC) is displayed in px^2.

There was a significant DilemmaXAge Group interaction effect on MAD ($\chi^2(4) = 11.844$, p = .019). Pairwise comparisons did not show any differences in MAD in each dilemma between the age groups. We then analysed the effect of Dilemma within each age group. There was no significant effect of Dilemma in children's MAD ($\chi^2(2) = 4.937$, p = .085) nor in adolescents' MAD ($\chi^2(2) = 5.564$, p = .062). Adults' MAD did not differ either in each Dilemma ($\chi^2(2) = 0.764$, p = .682).

A significant interaction between Dilemma and Age Group was observed for MDabove ($\chi^2(4) = 12.834$, p = .012). Pairwise comparisons did not show any differences in MDabove in each dilemma between the age groups. Within each age group, there was no significant effect of Dilemma on children's MDabove ($\chi^2(2) = 4.901$, p = .086). No significant effect of Dilemma on adolescents'($\chi^2(2) = 5.085$, p = .079) and adults' MDabove ($\chi^2(2) = 1.398$, p = .497) was observed either.

Furthermore, there was no significant DilemmaXAge Group interaction effect on AUC ($\chi^2(4) = 4.965$, p = .290). Type of Dilemma did not have a significant effect on AUC ($\chi^2(2) = 0.906$, p = .635), neither did Age Group ($\chi^2(2) = 2.404$, p = .301.

There was a significant DilemmaXAge Group interaction effect on xflips ($\chi^2(4) = 18.165$, p = .001) (Figure 3.31). Specifically, simple effects showed that it was children that had significantly more x-flips than adolescents (p = .003) and adults (p = .001) in the IR-NS dilemma. When we analysed each group separately, we found that children made significantly more flips (($\chi^2(2) = 17.008$, p <.001), specifically in the IR-NS dilemma compared to the ER-IR (p<.001) and compared to the ER-NS dilemma (p<.001). In comparison, adolescents did not differ in the number of flips in each dilemma ($\chi^2(2) = 1.227$, p = .541) and neither did adults (($\chi^2(2) = 0.055$, p = .973), suggesting that the overall differences observed were due to the children's x-flips.

Lastly, a significant DilemmaXAge Group interaction effect was observed on RTs ($\chi^2(4) = 16.694$, p = .002). Simple effects showed that children were significantly slower than adolescents (p = .008) and adults (p = .007) in the IR-NS dilemma. When we analyzed each group separately, we found that children were significantly slower ($\chi^2(2) = 11.728$, p = .003), specifically in the IR-NS dilemma compared to the

ER-IR (p = .003) and to the ER-NS dilemma (p = .005). Adolescents also showed marginal overall difference in RTs ($\chi^2(2) = 6.052$, p = .049), but no significant differences in pairwise comparisons. Adults did not differ significantly in their RTs in each dilemma ($\chi^2(2) = 2.221$, p = .329.

Figure 3.31

Mean number of x-flips per Age Group across in the IR-NS dilemma.

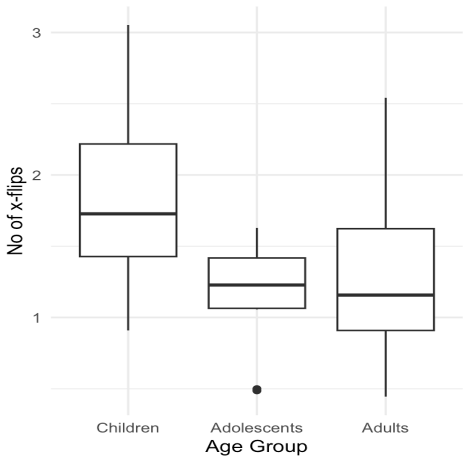

There was no significant DilemmaXAge Group interaction effect on the types of trajectories ($\chi^2(4) = 6.338$, p = .175). There was a significant main effect of Dilemma ($\chi^2(2) = 7.894$, p =.019), and specifically the IR-NS dilemma was significantly different from ER-NS (p = .019). To look further into these differences, we compared each type of trajectory per type of Dilemma separately (Table 3.14 and Figure 3.32). We found that the dilemmas did not differ significantly in frequency of straight ($F(2, 96) = 0.803$, p = .451), curved ($F(2, 96) = 1.595$, p = .208) or cCoM trajectory paths ($F(2, 96) = 0.401$, p = .671). There was also no significant main effect of Age Group ($\chi^2(2) = 2.095$, p = .351).

Table 3.15

Frequencies of trajectory types per Dilemma and Age Group.

	Straight	Curved	cCoM	dCoM
ER-IR	1236	238	160	0
ER-NS	1175	260	180	0
IR-NS	1247	220	155	1
Children	1371	302	203	0
Adolescents	771	161	122	1
Adults	1516	255	170	0

Figure 3.32

All observed trajectories as they have been automatically clustered in five different prototypes by mousetrap, a) Children, b) Adolescents, c) Adults

a)

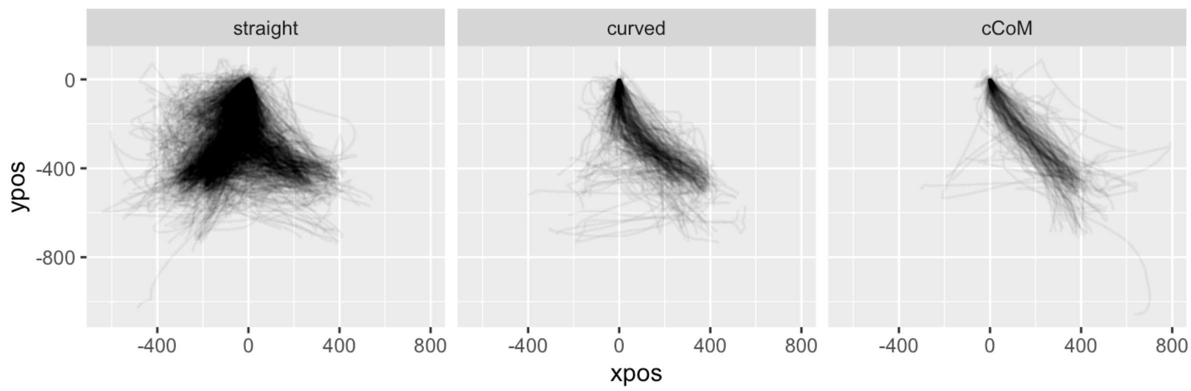

b)

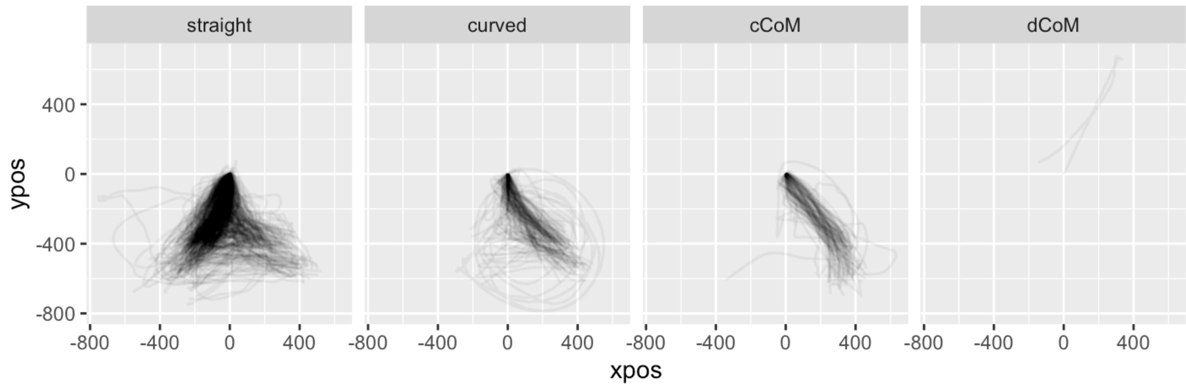

c)

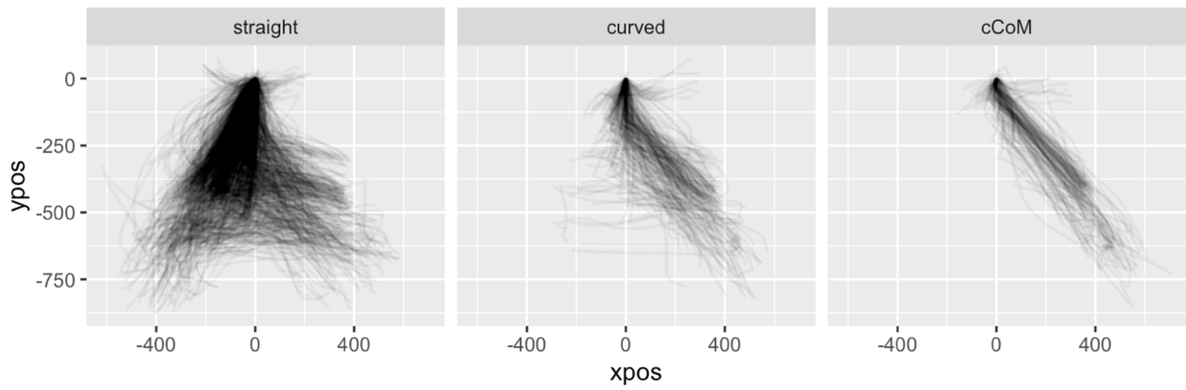

3.5.2.3. Secondary tasks

A subgroup of participants did not complete these tasks (2 adolescents and 1 adult), due to technical issues. No outliers were excluded.

3.5.2.3.1. Inhibition task

Task performance was measured as the number of errors in the No-Go trials divided by the total number of No-Go trials. We performed a one-way ANOVA with Age Group as factor (Figure 3.33a). There was no main effect of Age ($F(2,45) = 0.357$, $p = .702$).

Figure 3.33

Participants' error rate in the(a) Go/No-Go task and (b) Switching task.

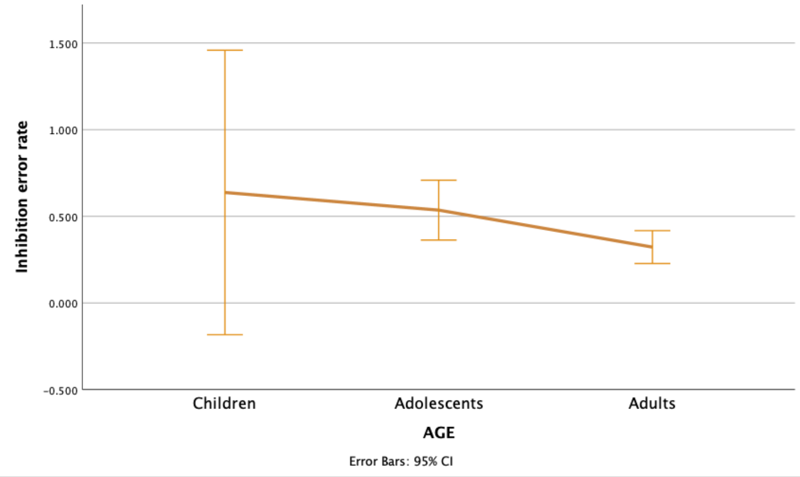

a)

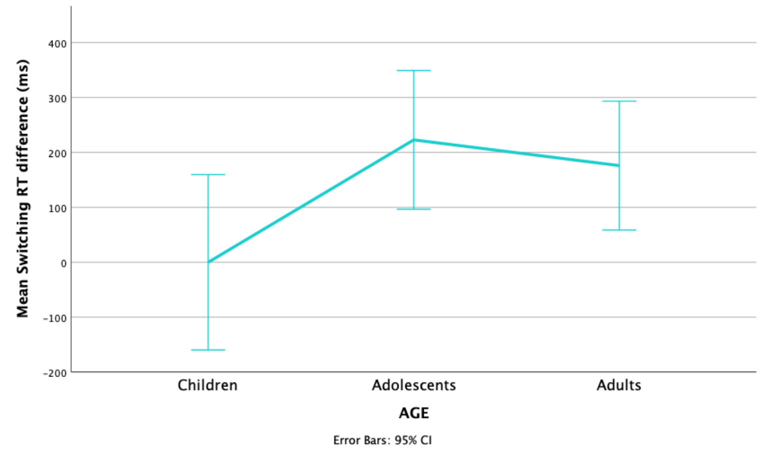

b)

3.5.2.3.2. Switching task

Task performance was measured as the difference between Reaction Times in mixed rule and simple rule trials. We performed a one-way ANOVA with Age Group as factor (Figure 3.33b. There was no main effect of Age ($F(2,45) = 2.623$, $p = .084$).

3.5.2.4. What predicts participants' choices?

We also tried to predict individual choices in the mouse-tracking task based on their scores in the secondary tasks. We built a Poisson generalised linear model entering Inhibition Score and Shifting Score as fixed effects.

In the full sample, the full model significantly predicted External Reward choices ($F(1,2) = 8.939$, $p = .011$). Specifically, Shifting Score was a significant predictor (.001 more ER choices, 95% CI, .000 to .001, $p = .020$). The model did not predict Informational Reward choices ($F(1,2) = 1.825$, $p = .402$), nor Novel Stimulation choices ($F(1,2) = 3.009$, $p = .222$).

We then ran analyses within each age group. In children, the model significantly predicted ER choices ($F(1,2) = 14.740$, $p<.001$). Specifically, Shifting Score was a significant predictor (.001 more ER choices, 95% CI, .000 to .001, $p = .001$). On the contrary, the model did not predict IR choices ($F(1,2) = 1.856$, $p = .395$), but it did predict NS choices ($F(1,2) = 14.930$, $p <.001$). Both Inhibition and Shifting Scores were significant predictors. Inhibition error rate positively predicted NS choices (.150 more NS choices, 95% CI, .069 to .232, $p<.001$), while shifting score negatively predicted NS choices (.001 less NS choices, 95% CI, -.002 to -.000, $p = .016$). To further examine the exact participants' preferences based on their EF scores, we first calculated the difference between ER, IR and NS choices in each dilemma. For example, the ERIR difference was a score measuring how many times participants chose ER over IR in the ER-IR dilemma, by subtracting the IR from the ER choices. Thus, a larger positive score indicated greater ER preference in the specific dilemma, and a negative score the opposite preference. We then correlated these scores with Shifting and Inhibition Scores. Inhibition error rate positively correlated with NSIR preference ($\rho = .592$, $p = .006$) but not with other differences, whereas Shifting score positively correlated with both ERIR preference ($\rho = .502$, $p = .024$) and ERNS preference ($\rho = .461$, $p = .041$).

In the adolescents group, the model did not predict ER choices $F(1,2) = 3.848$, $p=.146$), IR choices ($F(1,2) = 3.386$, $p=.184$), nor NS choices ($F(1,2) = 2.987$, $p=.225$).

Similarly, in the adults group, the model did not predict ER choices $F(1,2) = 1.421$, $p=.491$), IR choices ($F(1,2) = 2.785$, $p=.248$), nor NS choices ($F(1,2) = 1.878$, $p=.391$).

3.5.2.5. Summary of Results

In our fourth experiment, participants' performance differed from the previous findings. Specifically, participants in all age groups did not differ in their choices of the three options overall. However, a difference was observed in children's and adults' NS choices, with adults choosing NS more often. Otherwise, no difference was observed within each age group's choices either.

Our EF tasks did not show any age-related difference in performance, however when the scores were added in the regression model, they significantly predicted children's ER and NS choices. Specifically, higher switching performance was associated with more ER choices and less NS choices, while lower inhibition also predicted more NS choices, specifically in the IR-NS dilemma. The adolescents' and adults' choices were not successfully predicted by the model.

Analyses of the movement indices showed that children did significantly more x-flips and responded more slowly in the IR-NS dilemma, compared to the other groups and the other dilemmas. There were also significant interactions between type of dilemma and age group in the MAD and MDabove measures, but no significant differences in comparisons. The comparison of trajectory types showed that the IR-NS dilemma was significantly different to the others, but no specific type of trajectory appeared to be more common in this specific dilemma.

3.5.3. Discussion of Experiment 4

The aim of our fourth experiment was to compare participants' choices when the ER option was no longer a final reward acquired after following a certain strategy (i.e., a goal), but a trial-by-trial acquisition of either a candy or a coin – as a proxy for secondary rewards of everyday life. This change aimed in comparing the rewarding feeling after obtaining an external reward to the rewarding feeling of uncertainty resolution or novel stimulation.

Our results suggested that participants were equally driven by all options. By inspecting the data, we can observe a general tendency to follow similar patterns to the previous experiments, specifically for children and adolescents, whereas adults seemed completely divided between the options. This behaviour possibly highlights the importance and strong motivation which accompanies a goal especially for adults, regardless of the actual reward's value, which did not differ in Experiment 4 compared to the previous ones. However, in our first experiment adults seemed to balance their choices options too, despite showing some preference for the ER. In our second experiment, they still balanced their choices between the two options which included a goal (explicit and implicit), and chose NS less. However, the decision time-limit in Experiment 2 might have reinforced these goal-directed strategies even more, which could not have been the issue in the first and fourth experiments. As a result, we cannot safely conclude that including a goal changed adults' behaviour dramatically, although obvious differences are observed.

Furthermore, children and adolescents also seemed to be affected by the lack of an explicit goal. However, in all groups we observed large variations, suggesting that some individuals might have actually followed strategies more or less consistent to previous findings. Specifically in children, some of this variation seemed to be explained by their EF maturation – such that participants with a larger ability to switch (or, at least, with the ability to switch with less cognitive effort) chose ER more and NS less often. Moreover, participants with worse inhibition within the children group also chose the NS option more. The fact that this finding is observed in this experiment, but not the previous ones (where the ability to follow goals was more necessary) is intriguing. Especially regarding children's shifting score, we would expect it to predict a greater balance between the options, considering that it reflects the ability to successfully change rules in each context. However, it could also be the case that, even though many children might have wanted to choose the ER option, only the children with better ability to do this could keep this in mind when the type of dilemmas changed trial to trial, whereas the rest might be more easily sidetracked when faced with a different dilemma; e.g., an IR-NS one, where no ER option was provided, and found it harder to switch back to an ER driven approach. At the same time, the NS option seems to be one that children possibly tried to suppress, even when they could not – considering that children with lower inhibition and shifting ability chose it more. The explanation for this is not clear, as no instruction or external reason was provided for such a preference. They could have had such a behaviour towards any of the other options as well. We could speculate that this option possibly held greater risk, as it could not be predicted, and children's developing risk aversion influenced their tendencies.

The difficulty of the IR-NS dilemma for children is also clearly reflected in their movement trajectories, specifically in their movement entropy and time. This is consistent with the findings of our previous experiments, suggesting that children find the NS option tempting even when they end up choosing different options.

We will be discussing our overall findings, limitations and proposed changes and extensions in the General Discussion of this chapter.

3.6. Comparison across all four experiments

The analyses in each of our experiments revealed interesting findings regarding participants' choices in different age group and pointed towards more systematic differences in these different ages. However, in each experiment we chose to manipulate the decision environment, creating different dilemmas for groups, or we slightly changed the age limits. A summary of these differences is presented in Table 3.15. Based on these manipulations or differences, we considered interesting to examine whether participants' specific preferences are stable or whether the groups are significantly affected by facing different options or goals.

We expected that different decision contexts will influence participants' preference for each of the choices, especially the differences between having a goal or not – being instrumental to a goal will make an option more attractive. Furthermore, we expect that overall, children will consistently prefer the NS option significantly more than the other groups, and the adolescents will prefer the ER more.

Figure 3.34 shows each Age group's choices in all four experiments. We found a significant three-way interaction between Dilemma, Age Group and Experiment ($F(7.413, 352.134) = 3.060$, $p = .003$). Simple effects showed that different groups chose each option differently based on the experiment. Children and adolescents chose the ER option equally in all experiments, regardless of the presence of a goal (non-significant effects). However, adults chose the ER option less in Experiment 1 compared to Experiment 2 ($p = .008$). Furthermore, all groups chose the IR option less in Experiment 1 compared to the other experiments, as it was expected by the block termination when the ER was acquired. Similarly, children and adolescents also chose the NS less in Experiment 1 compared to the other experiments. However, adults chose the NS equally in Experiment 1 and 2, and significantly more in Experiment 4. Children also chose the NS significantly less in Experiment 4 compared to Experiment 2.

Table 3.16

Differences in age range and task design of the four experiments

	Age Groups	ER Goal	IR Goal	NS stimuli version	Decision Time limit
Exp 1	5-9 year olds 13-16 year-olds Adults	Yes	Yes	1	No
Exp 2	5-7 year olds 13-15 year-olds Adults	Yes	Yes	1	Yes
Exp 3	5-7 year olds	Yes	No	2	No
Exp 4	5-7 year olds 13-15 year-olds Adults	No	No	2	No

Figure 3.34

Choices per Age group in all experiments.

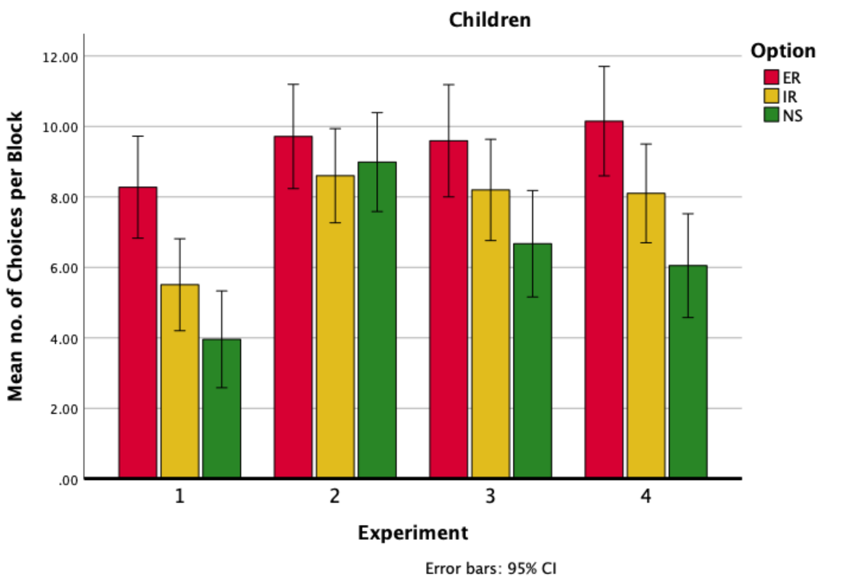

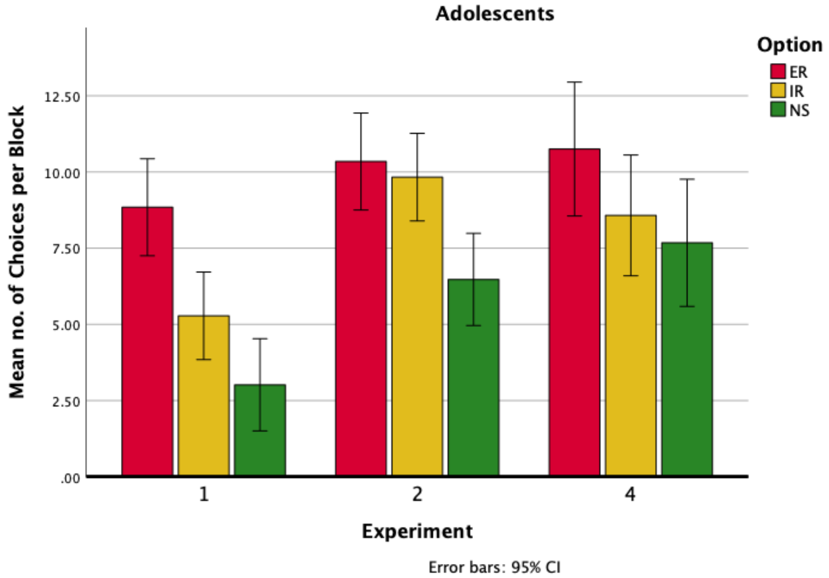

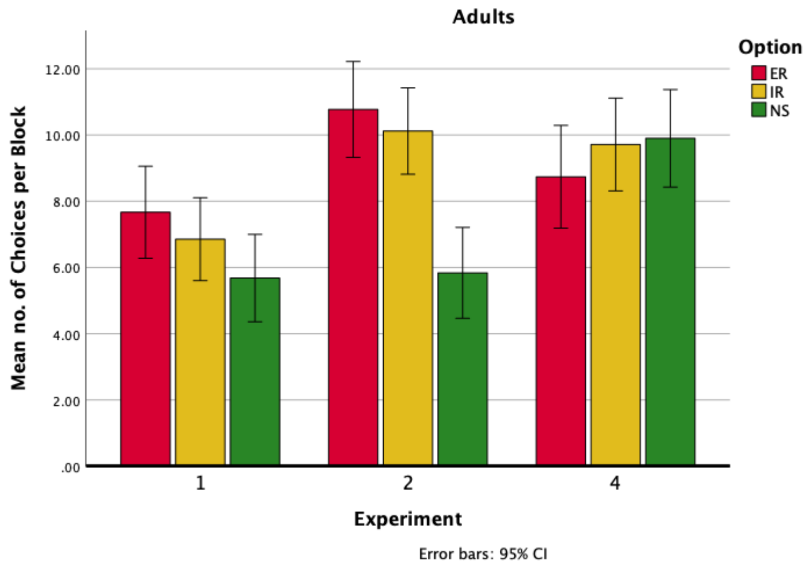

Taking all experiments into account, children and adolescents seem to prefer the ER option to the IR option (p = .004 for children and p = .023 for adolescents) and the IR option to the NS option (p = .031 for children and p = .002 for adolescents). In contrast, adults seem to choose the ER and IR options equally (p = .987) and prefer both options to the NS (p = .002 and p = .011 respectively). Overall between groups, the only significant difference seems to regard the IR option – adults prefer this option more compared to children (p = .041).

3.7. General Discussion

Our series of experiments aimed to uncover developmental differences in the balance between exploitation, exploration based on uncertainty, and exploration based on novelty. Based on previous findings, we started with the overarching hypothesis that children will be more novelty-driven than the other age groups, and that adolescents will be strongly exploitative, while we expected adults to be less novelty-driven and overall more exploitative than children. Furthermore, we were interested in the role that cognitive control and stable predispositions towards uncertainty might play for the individual differences in the explore/exploit balance. We found evidence supporting some of our initial hypotheses, but overall the findings paint a more complicated picture.

The expected children's increased preference for novelty was not consistently confirmed. Children did not differ from the other groups in their NS choices across all experiments, and only showed such a preference in Experiment 2. It seems that experimental manipulations played an important role in this difference. Specifically, children chose the NS option significantly more only when they were given a time limit for their decision. We consider these rapid decisions to be more revealing of children's initial, "intuitive" preferences, some of which they have to overcome in favour of more future-oriented rewards. Recent findings have shown that (adult) decision making related to the explore-exploit dilemma and behaviour towards uncertainty is affected by time pressure (Wu, Schulz, Pleskac, & Speekenbrink, 2022). It has previously been shown that imposing limitations to cognitive resources in decision tasks changes the participants' strategies to faster, "cheaper", more intuitive decisions (see also Kahneman, & Frederick, 2002), making immediate outcomes more salient (Ariely & Zakay, 2001). Specifically in an explore-exploit task, limiting cognitive capacities has often been shown to increase exploration, by leading to more risk-taking and making more decisions regardless of their actual outcome or expected value (Madan, Spetch, & Ludvig, 2015; Olschewski & Rieskamp, 2021). In contrast, it has also been shown to decrease exploration, by making people repeat their previous actions (Betsch, Haberstroh, & Molter, 2004). In our specific example, it is possible that putting cognitive constraints to children's decision-making influenced either their calculation of relative values (i.e., they often did not have enough time to analyse the possible extra future value they will gain by the ER and IR choices) or even after calculating the values, they had to resist/inhibit their NS-driven tendencies and choose more according to longer-term goals, on the action level – and this might require more cognitive power. Furthermore, since time-limited tasks lead to simpler strategies, children probably repeated the same strategy throughout the blocks – an assumption which is supported by their choices being consistent before and after the acquisition of the first reward. Such repetition of a simple strategy possibly explains the lack of conflict reflected in the mouse-tracking trajectories. While more difficult inhibition should be reflected in their ER-NS and IR-NS dilemmas, it is possible that participants chose early on not to do an exact comparison of the values each time, but repeat their first choices of each block, and thus no conflict was happening online in each trial. Indeed, in the other experiments (which did not have a time constraint), all participants – including children – found the IR-NS dilemma harder. The implication of cognitive control especially for children is further supported by our findings in Experiment 4, where participants with stronger EF skills were better able to inhibit the NS-driven and favour the ER-driven behaviour.

Our findings in the adolescents group validated our hypotheses that they will be very ER-driven, especially when the experiments included an ER goal. However, they were not significantly more ER-driven than the other groups in their total number of choices. Their behaviour was distinct from both the children and the adults' one in some interesting aspects. For example, in Experiment 2 they changed their strategy after acquiring the ER, showing a clear ER preference, and switching to an IR-driven one later on, similar to the adults' group, which still reflects an implicit learning goal. Their NS preferences are also similar to both groups in different contexts. Their difference with the children's NS choices are more obvious in the second experiment, where children clearly choose NS more – for reasons discussed above. However, adolescents do not seem to experience any conflict in these dilemmas (ER-NS and IR-NS), apart from the first experiment, possibly because the presence of a strong ER goal which could be acquired and lead to the termination of the block made the other two options equally (not) interesting for adolescents.

Similar observations were made for adults, especially regarding their ER-IR balance, although adults seemed to like both options equally, or they preferred a balanced strategy compared to a serial one (something possibly reflecting adolescents' strong desire/anxiety around the ER acquisition). Interestingly, adults seemed to choose the NS option a lot too, especially in the experiments where they could spend more time to calculate each value – while they seemed to focus almost exclusively in the other two options when the time was constrained. This behaviour is the opposite from the children's one, revealing that NS choices for adults were probably a reflection of "safe" exploration done with rationality when they had a longer time horizon, rather than impulsive and driven by compromised action suppression. However, it could be the case that repetition also affected adults' behaviours, as they started the block by choosing the options that were more valuable to them (ER and IR) and it was possibly more computationally efficient to keep choosing in the same way. Furthermore, the lack of an ER goal in Experiment 4 seemed to affect adults more than the other groups in their NS choices, possibly by making the ER less appealing and giving more choices to be spent exploring. However, the nature of the ER goal in Experiment 4 (candies/coins) might just not have been as successful as a proxy for a secondary reward in general. Although representations of money or sweets are often used as rewards in bandit tasks, the task instructions often inform the participants that they can win an amount of these money or candies at the end of the experiment. We did not give such an instruction, as we expected just the association with the actual objects would be sufficient.

Overall our methodology had several limitations. An important one regards the operationalisation of our concepts (external reward, informational reward, novelty, goals). While overall participants seemed to respond differently to the options they were given, suggesting some existing fundamental difference between them, the exact amount of uncertainty resolution/information gain (in the IR option) in each step in our first two experiments was not quantified. This was changed in the last two experiments, where we aimed for a specific amount of perceptual uncertainty, created with the same image resolution changes in all images. Still, the initial information held by each picture (in terms of complexity, symmetry, etc.) fluctuated randomly, and possibly influenced participants' interest/preference. Furthermore, due to the design of our task (2AFC), we could only analyse participants' exploration strategies by looking into broad preferences for each option, and not analyze their choices in a simultaneous setting. This would allow for a temporal perspective into decision making, and specifically to look into participants' profiles in each age group, when they might switch from each type of reward to another, and the contribution of learning, as it is eloquently shown in many recent paradigms of exploratory tasks in children and adults (e.g., Poli et al., 2022, Wilson et al., 2021).

The design of our secondary tasks also restricts the conclusions we can make based on our data. The utilisation of different EF tasks versions in our experiments limits the comparison between them and weakens our results. Furthermore, the uncertainty task in Experiment 2 seemed to particularly provide some interesting directions regarding the relationship between more stable attitudes and exploratory behaviour – and a longer version should be used with an explore/exploit paradigm to further investigate this relationship, as an alternative to commonly used questionnaires.

Finally, the inclusion of mouse-tracking methodology in such a paradigm offered useful and important complimentary findings. As these methods are rarely used in developmental value-based decision-making paradigms, our task is a good example of the processes that can be revealed, especially as to reveal underlying conflict and processing difficulty when values are compared (as it has already been shown in pure cognitive tasks; e.g., by Erb et al., 2016). Our findings also raise the possibility that children's conflict is revealed in different movement indices – children seemed to make more small but discrete changes such as x-flips and less curved trajectories which imply a continuous decision process. However, more detailed analysis on the data would be needed to reach conclusions on these topics. A more bottom-up process (e.g., as in Maldonado et al., 2019) would be needed to clarify whether children's decisions also share specific different characteristics in their trajectories compared to older groups. Apart from this question, our analyses could also be expanded by analysing the temporal characteristics of the trajectories, such as velocity and acceleration, but most importantly by breaking down the movements in theoretically important stages, in an attempt to identify the exact timescale of value-based decisions. This last approach, if combined with neuroimaging in such a paradigm, could possibly test the hypothesis about whether immediate reward and higher-order, future goals (both external and informational) influence movement serially or simultaneously (e.g., as other attributes in similar tasks; Sullivan et al., 2015).

In summary, our experiments provided some evidence for differential valuation and preference for external rewards, uncertainty resolution and novelty by school-aged children compared to adolescents and adults, and the possible implication of cognitive control maturation to these attitudes. The incorporation of mouse and finger tracking overall proved a promising technique to reveal children's preferences.

In our next chapter, we are examining how visual and haptic object complexity influences preschoolers interest and overall liking, attempting to dissociate the two processes and expand previous studies from vision to the haptic modality.

Chapter 4

Object exploration with vision and touch – Effects of object complexity on preschoolers' interest and aesthetic judgements[5]

4.1. Introduction

People casually perceive and characterise images, concepts or situations as simple or complex; such a characterisation is usually based on how much information must be processed. For example, musical compositions can be complex if they involve many changes in tune and melodies, or a political issue might be complex if many points of view are contradicting. Such judgments also take place in the perceptual domain, in a seemingly more automatic, immediate way – humans can judge whether an image or an object is simpler than another in a similar way as when judging objects' colour or size. The ability to track stimulus complexity spontaneously has been suggested to have biological significance, related to the extremely well-tuned human ability to identify learning possibilities in their environments; i.e., to track the available information and its potential for improving the observer's state of knowledge. Indeed, stimulus complexity has been long known to influence human information sampling and attention across the lifespan. Early theorists such as Berlyne (1960) have described and extensively investigated the effects of complexity on exploratory behaviour. In his arousal theory, Berlyne included complexity to the collative variables (i.e., the stimulus properties more likely to generate arousal) and suggested that medium amounts of complexity lead to increased motivation for exploration (curiosity) and more information-seeking behaviour. Similarly, Attneave (1954; 1957) took an information-theory approach to describe how specific variables of image complexity (e.g., recurrence or proximity) might lead to longer processing of visual stimuli. Since the days of these early studies, extended research in cognitive psychology with a focus on learning has

[5] This chapter is part of a larger collaborative project, with Tommaso Ghilardi, Francesco Poli and Sabine Hunnius, who had significant contribution to the conceptualisation and data analyses of this experiment.

shown that visual complexity is picked up even from infancy (Brennan, Ames, & Moore, 1966; Cohen, DeLoache, & Rissman, 1975), and that more complex sequences are also preferred by infants[6] (Addyman & Mareschal, 2013; Kidd et al., 2012). A corresponding but distinct line of research, experimental aesthetics, has also extensively investigated complexity preferences across the lifetime (e.g., Güçlütürk, Jacobs, & van Lier, 2016), promising more robust methods to quantify perceptual complexity, which has generally been overlooked by experimental psychologists. Comparing findings from the two fields gives rise to important questions and directions for further investigation.

One such question regards the nature of the psychological effects of stimulus complexity. As mentioned above, some early studies have related the arousal potential of a stimulus to its complexity with an inverted U-curve (e.g., Berlyne, 1970, Berlyne & Crozier, 1971; Kidd et al., 2012); when presented with a set of stimuli, participants will be more interested to the ones of medium complexity. However, others have failed to find evidence for such a relationship (e.g., Martindale, Moore, & Borkum, 1990). This inconsistency has been attributed to different factors. To start with, a dissociation between interestingness and liking/pleasantness of stimuli is necessary. While 'interestingness' is mostly associated with learning/categorising processes, quantifying the learning potential of a stimulus, 'liking' or 'pleasantness' reflects an aesthetic judgement and feelings of satisfaction and enjoyment. This dissociation is also likely to be reflected in the different measures that have been used in complexity preference studies. While most studies in experimental psychology considered looking time as a measure of preference in general (an assumption which still holds), it has been shown that looking or exploration time might specifically reflect stimulus processing in terms of learning, contrasted to explicit liking/beauty judgements (e.g., Wohlwill, 1968). Even in explicit judgements, when participants were asked to rate interestingness vs. pleasantness, interestingness seemed to increase linearly with increasing stimulus complexity or follow a U-shaped curve, while pleasantness peaked either in low or high degrees of

[6] From an information-theory perspective, a dissociation between complexity and novelty/unpredictability might not be meaningful. However, for the purposes of this study, we will refer to complexity as the amount and structural organization of information in the spatial domain. The two studies mentioned here focus on how information is organized in the temporal domain, closer to stimulus unpredictability.

stimulus complexity (W-shape, based on individual ratings; Berlyne, Ogilvie, & Parharn, 1968; Day, 1967;1968; Musinger & Kessen, 1964). Similar dissociation has been shown in school-aged children (Hutt & McGrew, 1969).

Regarding what might be causing this dissociation, it has been proposed that different informational aspects of stimuli might influence these two types of judgements. Indeed, complexity is better understood as multidimensional: one dimension involves the amount and variety of elements, while the second involves their organisation and grouping (Nadal, Munar, Marty, & Cela-Conde, 2010). Earlier studies also make such distinctions. For example, Chipman (1977; Chipman & Mendelson, 1979) distinguished between a quantitative factor (related to amount of elements) increasing complexity, and a structural factor (determined by different forms of structural organisation, but mainly by symmetry), which decreases complexity. When asked to provide ratings based on the images' asymmetry (a structural organization component), participants judged more asymmetrical images as more interesting but less pleasant (Day, 1968). More specifically, asymmetry has been shown to have differential effects on judgements of beauty (e.g., Jacobsen & Hofel, 2001). Consequently, research designs looking into the psychological effects of complexity should include both implicit and explicit measures, and carefully operationalise interest and liking. Furthermore, it is important to quantify complexity in terms of all the relevant dimensions, and carefully manipulate its changes.

However, even after dissociating these two aspects of preference, studies systematically report large individual differences in preference for complexity (Aitken, 1974; Güçlütürk et al., 2016; Lane, 1968; Rump, 1968), both in terms of interest and pleasantness. Specifically, Güçlütürk et al. (2013) used an exploratory clustering analysis to group people with low- and high-complexity preferences, and showed that different people tend to be consistent in their high or low complexity preferences. Moreover, they suggest that the frequently observed U-shaped curve can be the result of grouping all participants' responses together. Certain dispositional and contextual factors have been proposed as an explanation for individual complexity preferences. These primarily reflect more general approach-avoidance tendencies towards new information and uncertainty. In a recent application of emotion appraisal theories to the psychology of interest, Silvia (2005) proposed that feelings of interest presuppose both a potential for learning and a coping judgment; i.e., the belief that a particular learning goal can be achieved by the subject, taking into account their current total state (knowledge- and physiological-wise). Similarly, Gruber and Ranganath (2015) incorporate an appraisal step in their PACE (Prediction-Appraisal-Curiosity-Exploration) framework, suggesting that learners evaluate whether new information will lead to learning or harm (and this evaluation leads to exploration or avoidance).

While the appraisal is definitely contextual (e.g., influenced by the current mood), it is also affected by dispositional factors. For example, high scores in specific personality factors such as Openness to Experience and Contentiousness have been positively correlated with preference for more complex stimuli (Chamorro-Premuzic, Burke, Hsu, & Swami, 2010). However, when it comes to complexity specifically, these dispositions might also be confounded by experience: since complexity is a collative variable, it is always compared to how much experience an observer has with complex images. McDougall, Curry, and de Bruijn (1999) measured complexity and familiarity judgements of stimuli and found that less familiar shapes were also judged as more complex. This is an important point that should be considered when participants are exposed to similar stimuli over consecutive trials during studies. In summary, we could expect that individuals might have relatively stable dispositions towards complexity and certain exploratory strategies, which might consistently produce different responses in complexity preference tasks.

The aforementioned research has focused almost exclusively on the visual domain. However, complexity can also be relevant in or sensory modalities such as the haptic domain. The exploration in the visual and haptic domains have been shown to have many similarities (Lacey, Campbell, & Sathian, 2007; Lederman & Klatzky, 1987). Moreovoer, significant overlap has been shown in brain areas preferentially activated when assessing geometric representations regardless of their encoding modality (Amedi, Jacobson, Hendler, Malach, & Zohary, 2002). Thus, many common informational attributes are available and recognisable by both vision and touch, including both quantitative and structural information; e.g., variability of items or proximity/symmetry (Locher & Simmons, 1978; Overvliet, Krampe, & Wagemans, 2012; Overliet & Sayim, 2016).

A strong correlation in complexity judged through vision and touch has been reported in early studies (Owen & Brown, 1970). Although haptic exploration strategies in the presence or absence of vision have been well studied, very few studies have tried to measure the interestingness and pleasantness of haptic stimuli, in relation to their complexity. One study by Jakesch and Carbon (2012) showed increased liking of complex haptic stimuli after familiarisation (known as the *Mere Exposure* effect) based on the objects' texture (e.g., more liking for stone objects vs. wood). In a more recent study, Muth, Ebert, Markovic and Carbon (2018) had participants rate the pleasantness, liking and interest for more or less complex configurations of specific 3D objects (fish). They were particularly interested in identifying an insight effect on pleasantness (but not interest) through haptic exploration, similar to the one which can be observed through vision (i.e., images that cause sudden insight are rated as more pleasant). Indeed, increased complexity was strongly correlated with interest, while insight was correlated with both pleasantness and interest. Interestingly, the researchers made the same measurements for vision, compared the correlations between the two modalities, and found similar effects in both modalities.

In the current study, we aimed to examine how complexity might influence interest and liking through both vision and touch. We were further interested in how these preferences might develop from a very young age (i.e., preschoolers). From the scarce evidence mentioned above, children's preferences have only been measured in visual complexity studies, and the findings are often contradictory. Although haptic modality development and its relationship to object exploration and learning has been extensively studied (e.g., Pereira, James, Jones, & Smith, 2008), the effects of complexity have not been investigated. We were particularly interested in this age group because multisensory integration processes are still under development (Gori, Del Viva, Sandini, & Burr, 2008). This might affect children's ability to facilitate the mental representations of objects using mental imagery (i.e., crossmodal correspondences to unify how an object looks and how it feels when it is manipulated). As a result, testing children might enable a dissociation between liking as a result of direct perception vs. imagined manipulation to emerge, something that cannot be achieved in adults.

In our task, 4-year-olds explored 3D objects of three levels of complexity through vision and touch unimodally. The objects were based on abstract, symmetrical shapes (see Stimuli) used by Gartus and Leder (2017; constructed by Jacobsen and Höfel, 2002). Their study included both symmetric and asymmetric stimuli and aimed at comparing subjective (human) and machine-learning based ratings of complexity, revealing high correlations between the two. In our study, we only used symmetric stimuli, and manipulated complexity by changing the amount of variability on each shape (based on the number of different groups of characteristics; see Alberti & Witryol, 1990). Importantly, we did not modify the abstract patterns, but chose them based on complexity ratings in Gartus and Leder (2017; see Stimuli in Methods). We measured both exploration time (looking or touching) as a measure of interest and explicit ratings on liking and play preference.

We expected participants' exploration times to differ based on three possible strategies: (i) an overall preference for a specific level of complexity applying to all participants (e.g., more interest for the intermediate level), (ii) different preferences for individual participants (e.g., medium and high complexity explorers, consistent across modalities), and (iii) different preferences for each participant, which was also different between each modality (possibly due to their proficiency at exploring in one modality over the other). Finally, we expected participants' exploration times and explicit liking answers to differ. Specifically, we expected more individual differences in liking, such that they might consistently prefer low or high complexity objects.

4.2. Methods

4.2.1. Participants

Our sample consisted of 29 children (14 females, mean age: 4.5 years). All participants were neurotypical and had normal or corrected-to-normal vision. The participants were volunteers recruited though the Birkbeck Babylab database.

4.2.2. Design

All participants had to complete two experimental blocks, with 10 trials each. One block consisted of Visual trials (i.e., participants could see the stimuli but they could not touch them), while the other block consisted of Haptic trials (i.e., participants could only touch the stimuli but they could not see them.) The block order was randomised for each participant. In each Visual trial, the participant had to look for as long as they wanted to three objects of different levels of complexity (Low, Medium, High), presented simultaneously through three different holes on a board (position 1, 2 and 3; see Stimuli). In each trial, the level of object complexity presented in each position was randomised. At the end of the trial, participants had to answer an overall *Liking Preference* question (LP: "Which stone did you like more?") and a *Play Preference* question (PP: "Which stone would you play more with at the end of the game, if I were giving one to you?"). In each trial, the total exploration time (ET) for each object was measured, as well as their answers to the two questions above. Similarly, in each Haptic trial, participants had to touch three objects of three different levels of complexity with both hands, through side holes in three boxes, for as long as they wanted, and then answer the same questions. Total exploration time for each object and the participants' answers were recorded.

4.2.3. Stimuli

We used 18 objects of low, medium and high complexity as experimental stimuli (Figure 4.1a). The 2D versions (i.e., images viewed from the top, Figure 4.1b) of the objects had been objectively and subjectively rated for visual complexity in a previous study by Gartus and Leder (2017). In this study, the authors compared different objective complexity measures to identify the ones that most accurately predicted participants' subjective complexity ratings. They found that, from all twenty parameters evaluated, mirror symmetry and the root mean square contrast (RMS) of the images when saved as Graphics Interchange Format (GIF) were the best predictors (Gartus & Leder, 2017). Since we only used symmetric images, and thus mirror symmetry was not relevant, we used the RMSGIF measure to categorise the stimuli we used to test the children. After obtaining the full set of images from the authors, we clustered them in three sets based on their RMSGIF. The 18 images we eventually chose had large correlation between their objective and subjective ratings ($r(16) = 0.95$). The 18 chosen images were then 3D-printed and rated in terms of their haptic complexity by 10 adult volunteers. The participants were specifically asked to categorise the objects in low, medium and high complexity, and their ratings were highly consistent ($r(9) = .89$). The 3D printed objects were all red and black (black base 10cmX10cm, red shaped protrusion of 4cm height) and they included a small magnet inside the base, in order to be attached on boxes for the visual trials.

Figure44.1

a) 3D-printed object exploration task stimuli, low to high complexity (left to right), b) original images from Gartus and Leder (2017) study

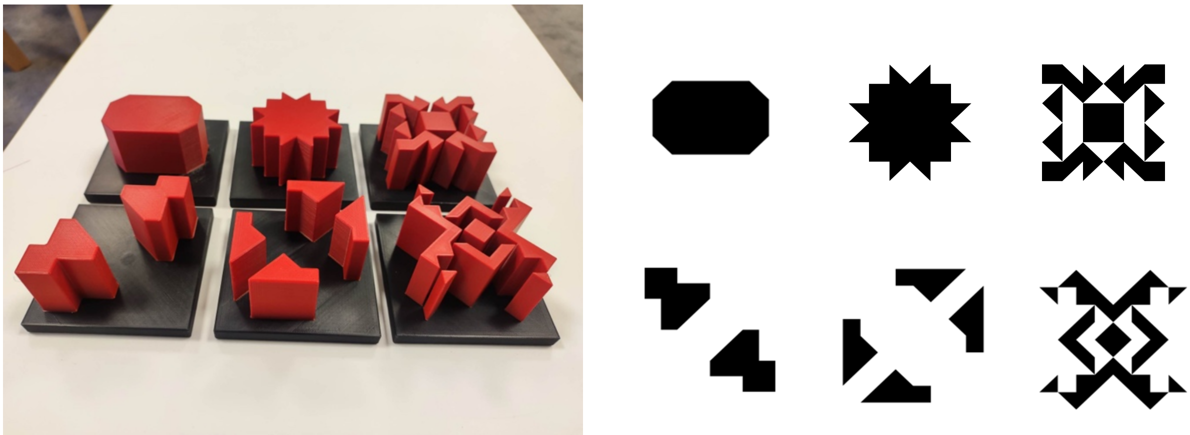

a) b)

We used three boxes with side holes and a removable top side to present the 3 stimuli in only 1 modality (Figure 4.2). For the visual trials, the boxes were positioned one next to the other, behind a black panel/occluder, which had three small holes for the participant to look through (Figure 4.3a and 4.3b). In each trial, the objects were attached in the interior back wall of each box, where a small magnet had also been placed. For the haptic trials, the same boxes were removed from the back of the panel, and were stacked one on top of the other, fitting on pre-designed recesses (Figure 4.2a and 4.2b). The top box was closed on top with the removable part. The side holes of the boxes allowed the participants to fit their hands inside and explore an object in each box.

Figure 4.2

Configuration of the boxes for Haptic trials, a) participant side, b) experimenter side

a) b)

Figure 4.3

Experimental setup for the visual trials, a) participant side, b) experimenter side.

a) b)

4.2.4. Procedure

Children were asked to stand in front of the board (for Visual trials) or in front of the block of boxes (for Haptic trials). Then the experimenter started narrating the following: “Today we are going to explore two magical kingdoms: a magical forest and an underwater kingdom. Which one do you want to explore first?”. Then the child answered and experimenter presented ten different cards to the child, from the kingdom they have chosen. The child was asked to name the characters they saw, as part of initial engagement. After all characters were named, the experimenter explained the following: “Each character has three magical stones in their house, and you can sneak-peak and look at the stones through these holes for as long as you like/you can put your hands through these holes and touch these magical stones for as long as you like”. Then the experimenter gave one of the objects to the child for familiarisation (this object was not used in the test trials). When the child was ready, they picked a card from the pile and the first trial begun. In each trial, the child started exploring from the position of their preference; they were allowed to go back to a previously explored position freely as many times as they liked. When they stopped exploring, they were asked the GL and the PP questions and were encouraged to give only one answer (one object of preference). Next, they were asked to pick another character card and the next trial started. After one block of trials ended, the configuration was changed by the experimenter to continue to the other modality block of trials. Meanwhile, the child was presented with 10 new character cards, and was asked to name those. The same procedure was followed for the next block.

4.2.5. Video coding

Two people coded the data independently – one of them was naive to the purposes of the experiment and the other was the experimenter. The dependent variables of interest were participants’ Exploration Time (ET), their overall Liking Preference (LP) and their Play Preference (PP). As participants were allowed to start exploring from the position of their choice, their Order of exploration was also coded.

Specifically, we coded as Exploration Time as follows: in visual trials, the total time in ms between the moment the participant turned their head towards a hole and aligned their eye to it, such that the object could be visible through it, and until they turned their head away from the hole. If the participant returned to look through a hole again, this time was added to the total time. As a result, Exploration Time refers to all the exploration attempts the participants did. In haptic trials, we coded as Exploration Time the total time in ms from the moment participants touched an object inside the box until they removed their fingers from it. Again, if they returned to the object, that time was added to the total time.

Participants' Liking Preference was the object they chose verbally or by pointing when asked "which one did you like more?". If the participants did not seem sure (i.e., pointing to more than one object), they were asked again to give a definitive answer and this was the one we coded as chosen. The initial coding referred to the position of choice, which was then matched to the complexity level of the object of choice, as well as the specific object the participants were choosing. Similarly, their Play Preference was the object they chose when asked "which one would you play more with at the end?". Their Order of Exploration referred to whether they explored each object first, second or third.

4.3. Results

Three participants were excluded from the sample as they completed fewer than half of the trials in at least one Modality block. Therefore, the final sample consisted of 26 participants.

Table 4.1 shows the main descriptive values for Exploration Time. We used R (R Core Team, 2020) and *lme4* (Bates, Maechler & Bolker, 2012) to perform a mixed effects analysis of the relationship between ET and our predictors (Modality and Complexity). We also included Order as a covariate. We used the log-transformed values for RTs, as the raw scores violated the normality assumption. As fixed effects, we entered Modality, Complexity, Order and Trial, and the interaction between Modality and Complexity into the model. To incorporate the dependency among observations of the same subject and the repeated trials, as random effects, we had intercepts for subjects. The model for ET was the following:

ET.model = log(ET) ~ Modality*Complexity + Order + Trial + (1 |Subject)

P-values were obtained by restricted maximum likelihood ratio tests of the full model with the effect in question against the model without the effect in question (chi-square tests for nested models).

Table 4.1

Exploration Time means and SDs per Complexity level and Modality

	Visual			**Haptic**		
	Order 1	Order 2	Order 3	Order 1	Order 2	Order 3
Complexity 1	3043.29 (2740.99)	2202.32 (2248.95)	2255.27 (1658.76)	4287.44 (4685.67)	4192.38 (5116.28)	3525.59 (2628.35)
Complexity 2	3721.98 (3514.85)	2749.05 (3730.37)	2742.22 (2558.47)	6090.71 (6684.11)	3239.88 (3015.15)	3827.76 (3752.50)
Complexity 3	3884.93 (3935.78)	2523.74 (3207.97)	2575.17 (2591.29)	4899.40 (6113.13)	3273.00 (2802.75)	3836.93 (4029.92)

There was no significant interaction between Modality and Complexity ($\chi^2(1) = 1.893$, $p = .388$). No main effect of Complexity was found ($\chi^2(2) = 3.061$, $p = .216$). There was a main effect of Modality ($\chi^2(1) = 80.895$, $p < .001$), showing that participants spent significantly more time exploring in Haptic compared to Visual trials. There was also a main effect of Order ($\chi^2(2) = 61.29$, $p < .001$). Specifically, participants explored the first object they saw or touched significantly more than the others ($p < .001$), whereas the ET in the other two positions did not differ. Finally, Trial was a significant covariate ($\chi^2(1) = 214.71$, $p < .001$), negatively correlated with ET (i.e., participants explored significantly more time in the first three trials and significantly less in the final three trials). We used AIC model selection to distinguish among the aforementioned set of models, as well as the model without the random effect factor. The best-fit model, carrying 84% of the cumulative model weight, included Modality, Complexity, Order and Trial, no interactions and the random effects. Since we had specific hypotheses for the two modalities, we also analysed the data of each modality block separately. As the relationship between Complexity and ET is not always linear, we ran a polynomial regression for the Complexity predictor, in order to also check for a quadratic fit. A main effect of Complexity as a linear predictor was found in the visual block ($\chi^2(2) = 2.136$, $p = .033$). In the haptic block, Complexity was not significant as a linear predictor ($\chi^2(2) = 0.509$, $p = .611$) nor as a quadratic one ($\chi^2(2) = 1.597$, $p = .111$).

We also decided to conduct the analyses with Complexity as a continuous factor, by using the corresponding RMSGIF measure for each object. We did this in order to potentially capture more subtle influences of complexity which might have been missed when objects were grouped into three levels of complexity. There was no significant interaction between Modality and Complexity ($\chi^2(1) = 1.683$, $p = .194$). No main effect of Complexity was found ($\chi^2(2) = 0.356$, $p = .551$). There was a main effect of Modality ($\chi^2(1) = 82.965$, $p < .001$), as participants spent significantly more time exploring in Haptic compared to Visual trials. There was also a main effect of Order ($\chi^2(2) = 60.51$, $p < .001$). Specifically, participants explored the first object they saw or touched significantly more than the other ($p < .001$), whereas the ET in the other two positions did not differ. Finally, Trial was a significant covariate ($\chi^2(2) = 201.71$, $p < .001$), negatively correlated with ET (i.e., participants explored for less time as trials progressed). We used AIC model selection to distinguish among the aforementioned set of models, as well as the model without the random effect factor. The best-fit model, carrying 71% of the cumulative model weight, included Modality, Complexity, Order and Trial, no interactions and the random effects by individual subjects. Similarly, we also analyzed the data of each modality block separately. A main effect of Complexity as a linear predictor was found in the visual block ($\chi^2(2) = 5.011$, $p = .037$). In the haptic block Complexity was not significant as a linear predictor ($\chi^2(2) = 0.501$, $p = .617$) nor as a quadratic one ($\chi^2(2) = 1.882$, $p = .060$). Figure 4.4 shows the relationship between object complexity and ET in each modality.

We then analysed the relationship between the ET and the explicit object preferences (i.e., the answers to the Liking/Play questions). Table 4.2 shows participants' preferences in total, and in each modality. Complexity did not seem to significantly affect participants' Liking Preference ($F(2,52) = .001$, $p = .999$), nor Play Preference ($F(2,52) = .918$, $p = .406$) when comparing the data of both modalities together. When comparing participants' preferences in each modality separately, in visual trials Complexity did not affect participants' Liking Preference (($F(2,52) = .038$, $p = .963$), nor Play Preference ($F(2,52) = .034$, $p = .967$). Similarly in haptic trials, Complexity did not affect participants' Liking Preference (($F(2,52) = .047$, $p = .954$), nor Play Preference ($F(2,52) = 2.101$, $p = .133$).

Figure 4.4

Effect plot showing the relationship between a) Exploration Time and visual RMSGIF complexity, b) Exploration Time and haptic RMSGIF complexity

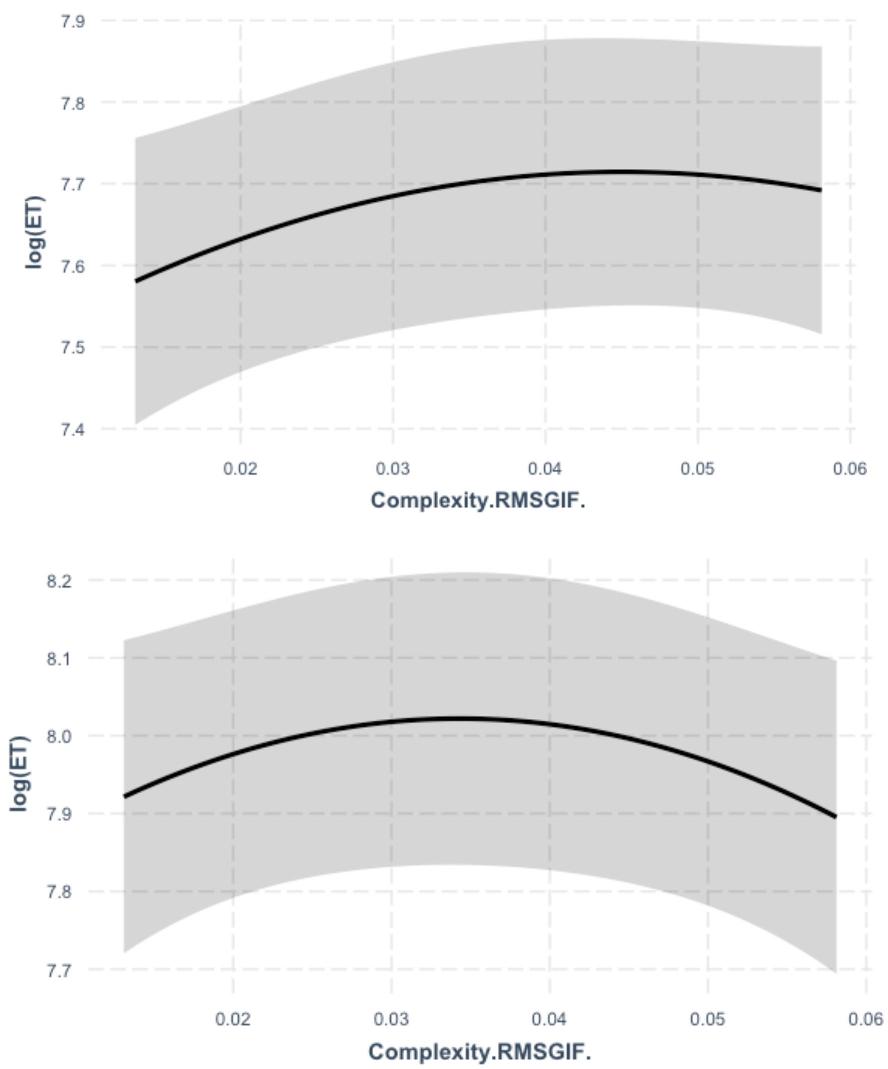

Table 4.2

Participants' mean frequencies of Liking and Play preferences per Modality and level of Complexity

		Low Complexity	Medium Complexity	High Complexity
Visual	Liking	9.78 (5.66)	10.22 (4.26)	9.89 (5.18)
	Play	9.78 (5.17)	9.67 (4.80)	10.11 (6.07)
Haptic	Liking	9.56 (4.71)	9.11 (4.74)	9.33 (4.35)
	Play	7.56 (4.02)	10.44 (5.29)	9.67 (4.80)
Overall	Liking	19.33 (8.25)	19.33 (6.83)	19.22 (7.07)
	Play	17.33 (7.46)	20.11 (5.81)	19.78 (7.45)

However, a visual inspection reveals very different patterns both between participants and between the two types of explicit preference in each modality. Some participants seem to prefer the most complex objects visually but not haptically (e.g., Participant 2; Figure 4.5c), while others' preferences might match in the Liking question and the Play Preference question across modalities but differ between them (e.g., Participants 1, 3, Figure 4.5a,b). Some might choose similarly in both questions in one modality but differ in the other (e.g., Participant 8, Fig 4.5d). As a result, we decided to cluster participants based on their preferences in each modality, and then use the clusters to predict their ET. We used the k-means clustering method to create a predetermined number of clusters for each modality and preference. To identify the optimal number of clusters, we compared different groupings using the silhouette method (i.e., the optimal number of clusters was the one with larger silhouette values). For Visual Liking, the optimal number of clusters was 3, corresponding to participants' larger preference for the Low, Medium or High complexity. For Visual Play preference, the optimal number of clusters was 6, corresponding to the aforementioned strong preferences (Low, Medium, High), as well as three categories for Low-Medium, Medium-High and Balanced preferences. Similarly, in the Haptic Liking, the optimal number of clusters was 6, which corresponded to Low-Medium, Medium-High, Balanced preferences, as well as Low-High and Strong Medium preferences. Finally, for Haptic Play, the optimal number of clusters was 6, corresponding to the same categories as in Haptic Liking. However, participants' clusters did not significantly predict their ET for either of the explored objects' complexity level in any modality (Appendix B).

Figure 4.5

Participants' mean frequencies of preference per level of complexity. Columns, from left to right: Visual Liking Preference, Haptic Liking Preference, Visual Play Preference, Haptic Play Preference. Rows show different participants, a) participant 1, b) participant 3, c) participant 2, d) participant 8

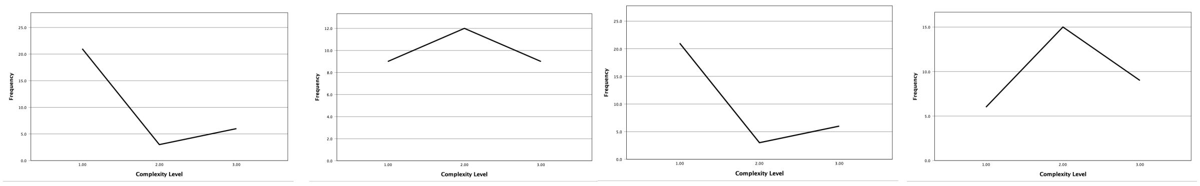

a)

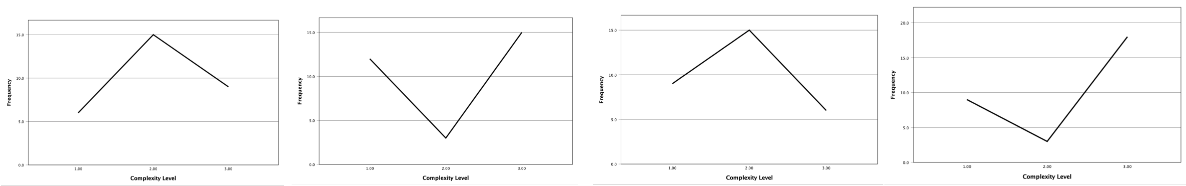

b)

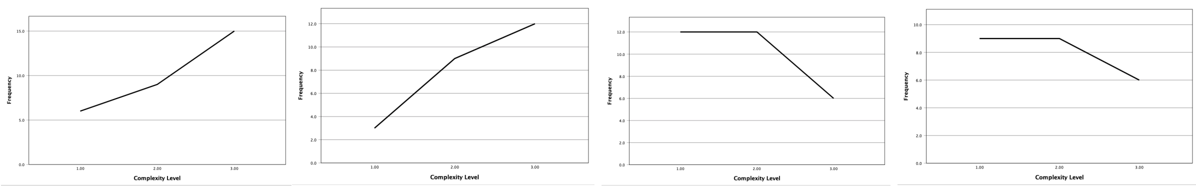

c)

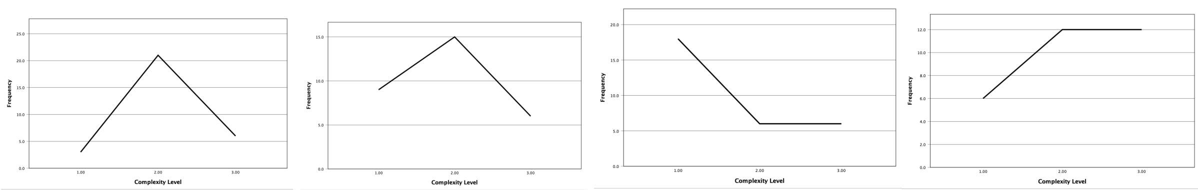

d)

4.3.1. Summary of results

Participants' exploration time was significantly influenced by the order they chose to explore the options, spending more time in the first position, as well as by trial number, and spending less time exploring as the task progressed. Furthermore, they spent more time exploring objects haptically than visually. There was no main effect of Complexity (as a category – Low, Medium, High – or as an exact measure (RMSGIF)) when analysing visual and haptic trials together. However, there was a main effect when the two modalities were analysed separately. Complexity significantly affected ET in the visual block, with ET increasing linearly with higher object complexity. In contrast, Complexity did not affect ET in the haptic trials, despite approaching a significant quadratic relationship. Participants' explicit preferences were not influenced by object complexity in the full sample. However, large individual differences were observed, and individual clusters revealed preferences for different levels of complexity. These preferences were also not consistent across modalities for all subjects. Finally, participants' ET could not be predicted by their explicit preferences, as grouped in separate clusters.

4.4. Discussion

Our experiment aimed to understand the effect of visual and haptic object complexity on preschoolers' interest and pleasure judgements. Specifically, we were interested in dissociating the two types of preferences, such that interest would be related to a learning process (reflected in participants' exploration time of the objects), while pleasure judgements would be explicit statements of preference, either regarding participants' overall liking of an object, or their hypothesized interaction with it. We expected the two measures to be affected by object complexity, and not to correlate with each other. Furthermore, we were interested in the possible differences of both of these preferences between the visual and the haptic domain. Participants could be consistent or differ across modalities – we refrained from making specific hypotheses here due to the exploratory nature of this study on the effect of complexity on the haptic domain. We found evidence for some of our initial hypotheses, while some of our findings need further investigation to be clarified.

Complexity was found to influence participants' exploration time in the visual domain, increasing it in a linear fashion. This suggests that participants could keep track of the information available and successfully process it, devoting the time needed to understand the structure of each object's surface. This linear relationship between visual complexity and time has previously been documented (Berlyne, 1963; Berlyne & Lawrence, 1964; Locher & Nodine, 1978). In contrast, we did not observe any significant relationship between haptic complexity and exploration time, although the effect resembled an inverted U-shaped curve, approaching a quadratic relationship. A linear relationship between haptic complexity and scanning time has been documented in an older study with adults (Locher & Simmons, 1978). We suspect that, in our case, haptic proficiency likely influenced participants' time spent on touching objects, such that the high-complexity ones were too difficult for preschoolers to process in detail, while the simple ones were too easy. Furthermore, the not fully developed correspondences between vision and touch might have made the processing even harder, as children could not form the image of the very complex objects assisted by vision (Gori et al., 2008). That said, other factors might have also influenced performance in the haptic trials; these will be discussed in more detail in the study limitation section below.

Participants' explicit judgements were not consistently influenced by object complexity in either of the two modalities (nor in total). The results of the overall liking question can be directly compared to previous findings (as these are mostly in the visual domain) but is also relevant to the haptic behaviour. Various studies have attempted to relate subjective visual complexity and aesthetic judgements (Aitken, 1974; Day, 1967; Güçlütürk et al., 2016; Nadal et al., 2010; Nath, Brändle, Schulz, Dayan, & Brielmann, 2023; Sun & Firestone, 2022), revealing high inconsistency in their findings. The seminal suggestion of an inverted U-shaped relationship (i.e., intermediate complexity items considered more beautiful) by Berlyne (1971) has been replicated by some (e.g., Lakhal, Darmon, Bouchaud, & Benzaquen, 2020) but others have documented a linear relationship (e.g., Day, 1967) or no relationship at all (e.g., Messinger, 1998). Güçlütürk et al., (2016) showed that the inverted U-shaped curve results from averaging individual responses, but does not explain participants' individual preferences. Indeed, our findings in the visual domain show large individual differences, such that participants are either low-, medium-, or high-complexity driven; this is even more the case in the haptic domain. Recent studies

(Nadal et al., 2010; Nath et al., 2023) have attempted to explain such differences by attributing them to three factors: (i) differences in how studies define, measure and manipulate objective complexity, (ii) diverse stimuli (often handcrafted) and (iii) large individual differences in beauty preferences, likely attributable to other factors (e.g., exposure/expertise). Regarding its definition and manipulation, complexity is a construct comprised of different aspects (number and variety of elements, organization of elements, symmetry). In our study, we only manipulated objects' number, variety and organisation of elements, and only used symmetric objects. However beauty judgements are specifically influenced by symmetry (e.g., Day, 1968, Eisenman & Gellens, 1968), so possibly the elements we manipulated tapped more into the amount of information and less into pleasantness – or at least less consistently.

Furthermore, we crafted our objects based on one of the proposed methods for measuring visual complexity (i.e., image compressibility; Marin & Leder, 2013), while recent studies have suggested possibly using more general measures. For example, Nath et al. (2023), used an algorithmic method (cellular automata) to generate visual patterns by systematically arranging squares and measured various informational aspects (e.g., entropy, symmetry, density), and dissociated beauty from complexity – here, subjective beauty was negatively predicted by disorder (asymmetry and entropy). Furthermore, a major limitation of our study was using a visual complexity measure to create haptic stimuli. Although we piloted the objects' judged complexity on some participants and they gave consistent answers, the exact manipulation during testing could not be controlled – participants also haptically scanned the sides of the objects, processing uncontrolled information. In their recent studies, Sun and Firestone (2021, 2022) used a method to create shapes that could possibly also be used to measure haptic complexity. They used a computational geometry approach to extract the skeletal complexity of shapes, which corresponds to a specific configuration including variety, organisation and disorder of objects. As this skeleton also applies to 3D objects, it could possibly influence manipulation in a more predictable way.

Importantly, we observed differences between participants' general liking and play preferences, suggesting that children had the ability, at least partly, and at least some of them, to imagine manipulating the objects, and the possible haptic and proprioceptive information they could get by doing it – dissociating this information from the visual ones.

Consistent with our hypothesis, participants' explicit preferences did not predict their exploration times, supporting a dissociation between interest/learning possibilities and pleasantness which has been proposed before (e.g., Day, 1967, 1968). This was also the case for play preferences, which could hypothetically predict haptic exploration time – as this involved manipulating the object. However, the lack of visual input likely made this type of exploration less appealing – in naturalistic play vision and touch rarely happen in isolation.

Some limitations of our study involve possible memory components which might have influenced especially the answers in the explicit judgment questions. Participants were allowed to explore positions previously visited if they wanted to, but very often they did not. This might have created recency effects in their answers – although the data suggested that the order of exploration played a role in the opposite way: participants seemed to explore the first positions for significantly more time than the following two.

Future modifications should be made to clarify our findings. The objects should be redesigned to better manipulate complexity differences in the haptic domain, and possibly include manipulation of object symmetry, to check for specific interactions between symmetry and amount/organization of elements on exploration time and judged pleasantness. Furthermore, participants' proficiency in visual and haptic discrimination should be tested with a relevant task.

In summary, our study compared the effects of visual and haptic object complexity on exploration time and pleasantness, in preschoolers for the first time, thereby extending existing adult research. Our findings support a dissociation between interest and exploration on the one hand and subjective beauty on the other hand, and introduce a more quantitative approach to manipulate object complexity, previously overlooked in developmental studies that have tended to focus on learning and attention.

Chapter 5

General Discussion

5.1. Overview of empirical findings

5.1.1. Object complexity influences the planning and execution of object-fitting actions in preschoolers

In Chapter 2, we investigated the role of object complexity in preschoolers' actual and imagined object manipulation, using an object-fitting paradigm. Our findings suggested that complexity influences the object-fitting process and identified possible uses of the hands in directing attention during spatial imagery. Based on previous studies on the development of object fitting (e.g., Fragazsy et al., 2015; Ornkloo et al., 2007), we expected that manipulating object complexity would directly influence children's fitting process and accuracy. We investigated an older age range than most previous studies because we were specifically interested in the process. As a result, almost all preschoolers were accurate in fitting our objects, but their fitting attempts increased in number with more complex objects. However, as our objects' complexity was not varied consistently across the symmetry and number/organisation of elements factors, we cannot safely conclude which specific factor influenced children's increasing difficulty.

Furthermore, we aimed to expand previous findings by examining the effects of object complexity on preschoolers' spatial imagery; i.e., whether more complex objects would decrease accuracy when children had to mentally rotate objects, and whether they would use embodied strategies (i.e., representational gestures) to support their mental processing. Such strategies have been previously documented in adults (e.g., Chu & Kita, 2011), while using pointing gestures to support memory has been previously shown in children (Delgado et al., 2014). Our findings showed that children did indeed use pointing gestures in a spontaneous fashion when they have to mentally rotate objects. They also seemed to use their fingers as a stable perceptual anchor in the environment to help them minimise information-gathering eye movements. However, this behaviour was not related to children's visual memory abilities and did not result in higher accuracy (children were equally accurate when they were not allowed to use their hands). It was also not related to object complexity – children seemed to use such strategies equally for all objects.

Finally, while we did not identify age differences in our sample regarding their accuracy, we did observe that participants who planned their fitting movement for longer (i.e., they used their eyes to gather information – and possibly imagine the object manipulation), were more accurate in their subsequent fitting.

Overall, these results point towards a flexible system for object-fitting in the preschool years. Children are able to keep track of objects' spatial information and use both mental and embodied resources to accurately guide the process of manipulating and fitting them.

5.1.2. Exploratory decisions vary across development – time pressure and cognitive control influence the balance between exploration and exploitation

In Chapter 3, we conducted a series of four experiments to investigate how people engage in exploitation, uncertainty-driven exploration and novelty-based exploration across development. We compared school-aged children, adolescents and adults in various versions of a decision-making task, while tracking their mouse positions to measure their ongoing decision process. Based on previous findings, children were assumed to be more exploratory than adolescents and adults (Somerville et al., 2017; Wu et al., 2019), and even further, to differ in the exploratory strategies that they engage in; e.g., children have been shown to engage in more random and more directed (or uncertainty-based) exploration than older participants. However, the developmental trajectory of novelty-based exploration is still unclear so we aimed in clarifying the separate contribution of uncertainty and novelty in exploratory choices. We were also interested in explaining why changes in this balance across development might occur. To this end, we targeted the maturation of cognitive control as a possible mechanism which might underlie the shift from exploration to exploitation more generally, as well as the specific preference for novelty-based compared to uncertainty-based exploration (Gopnik et al., 2017). Our overall findings suggest that children are not more novelty-driven or uncertainty-driven than adolescents and adults when allowed to choose in their own time. However, when time pressure is applied to the decision process, all the different groups' behaviours are affected, such that children are significantly more novelty-driven than the other groups, while adolescents are strongly driven by the external reward. As applying time pressure is related to limiting cognitive resources and leading participants to faster, more intuitive decisions (Kahneman & Frederick, 2002), this suggests that, when fewer cognitive resources are available, children cannot inhibit their action plans towards the novel option, or they cannot successfully calculate the values of options. Furthermore, all groups persevered more with their choices in the time-limited version of the experiment; this is also consistent with relying on a less cognitively-taxing strategy. The fact that cognitive control plays a role in these choices is also supported by the findings of our fourth experiment, showing that participants with better EF performance favoured the exploitative option compared to the novelty-based one.

The mouse-tracking analyses were expected to reveal greater decision conflict between equally valued options. Such conflict was revealed between the two exploratory options for all groups in most of the experiments but, in contrast to previous findings (e.g., Scherbaum & Kieslich, 2018), no conflict was evident when time pressure was applied. However,more variability was observed in children's movements during greater conflict choices, suggesting that conflict might manifest differently in the younger groups.

Overall, this series of experiments suggests that (i) novelty is an independent causal factor in exploratory choices, and (ii) cognitive control should be further investigated as a causal factor involved in value-based decisions across development. Furthermore, child-friendly versions of decision-making paradigms accompanied with mouse or finger-tracking can be developed to investigate how young participants calculate and compare values, an aspect that is often overlooked in similar tasks.

5.1.3. Visual and haptic object complexity differentially affects preschoolers' interest and pleasantness judgments

In Chapter 4, we investigated how objects' visual and haptic complexity influenced preschoolers' interest and explicit preference judgements, and the relationship between the two types of preferences. Our findings suggest that children can track changes in complexity and their processing is modulated by it, especially in the visual domain. Furthermore, we found no relationship between exploration time and explicit liking judgements, in either sensory domain. Previous research has reported contradictory evidence regarding the relationship between exploration time and visual complexity (e.g., Berlyne, 1967; Day, 1968), as well as between explicit beauty judgements and visual complexity (Aitken, 1974; Güçlütürk et al., 2016; Nadal et al, 2010). However, studies in the haptic domain are very limited. Moreover, most studies focus on adults' aesthetic judgements, while developmental studies tend to focus on the "interest" (related to learning) aspect of exploration, not on the liking/pleasantness aspects.

We measured exploration time as a proxy for participants' interest/learning value, showing that visual complexity linearly increases this time, while haptic complexity did not have a significant effect (although it aproached a quadratic relationship). In contrast, explicit liking judgments were not affected consistently by complexity, but rather, were made by participants based on their individual preferences. This adds to the existing literature by providing more evidence that the commonly found linear or inverted U-curve relationship between complexity and attention or preference is not consistent. Possible explanations of this difference in findings include both individual predispositions and sensory experience factors, as well as methodological ones. This case is even more apparent in the preschoolers' data, who are still experiencing sensory development (especially in the haptic domain and visual-haptic multisensory integration). Furthermore, different studies manipulate complexity by varying different factors (e.g., symmetry, organization of elements, see Nadal et al., 2010), which makes comparisons and conclusions hard. In our study, we did not manipulate image symmetry, which may have influenced participants explicit liking more robustly.

Overall, our study suggests that (i) interest and explicit preference should be dissociated when measuring the effects of complexity and (ii) preschoolers can track visual and haptic complexity in objects, but are constrained by their proficiency in each modality to identify the most informative ones.

5.2. Theoretical contributions

5.2.1. Increased novelty preference in children compared to older groups – the possible role of cognitive control

Older studies on the effect of novelty on attention and exploration have explicitly compared familiar to new stimuli (e.g., toys), while more recent examples focusing on learning have measured participants' ability to track the informativity of different options, comparing learnable (or unlearnable) unpredictable stimuli to novel stimuli. These more recent paradigms usually have a task goal (e.g., maximising rewards) and all stimuli are related to this goal. From this perspective, novelty signifies an informational attribute and the experimental question relates to whether participants consider it more or less useful in their quest to achieve the final goal. Compared to this approach, our studies conceptualise both uncertainty and novelty differently: apart from external rewards, the motivation is either perceptual uncertainty resolution for specific stimuli or novel stimulation in a much broader sense – and, importantly, *unrelated to a main external goal*. As a result, novelty preference in our paradigm is closer to the concepts of sensation-seeking, diversive curiosity or broad interest. However, since our tasks included mutually exclusive choices, they also bare similarities to the explore-exploit dilemma bandit problems – participants needed to sacrifice tokens if they wanted to explore. This implies that our results should be interpreted as dilemmas between exploitation and exploration, where exploration is either specific uncertainty ("information gap" closing) or sensation-seeking driven – without holding strong views about the continuous or discrete nature of such motives.

Across all four experiments in Chapter 3, children explicitly preferred acquiring the external reward and considered uncertainty resolution more valuable as compared to novelty. This is consistent with previous findings in this younger age group, which show that directed exploration is extensive in school-aged children – proof of their increasing ability to identify sources for information gain and to direct their exploration towards these. However, our experiments showed that novelty is still tempting for children. This was especially evident in their choices when they had time constraints – i.e., when they had to choose intuitively, and in the choices of participants with better vs. worse EF skills. The fact that such preference for novelty still exists in school-aged children suggests a dissociation between the importance of the two informational attributes in younger ages, such that novelty is still highly valued for learning progress. However, since our novelty option did not help resolve any task-specific informational uncertainty, it is more likely that the novelty preference reflects children's broader exploration tendency – the generation of learning possibilities which can be resolved is highly valued and adaptive in younger ages, as proposed by Gopnik's (2015). Based on the same theory, this high valuation can also be a by-product of immature cognitive control. It could possibly result from an underdeveloped ability to track the exact informational value in the environment, and thus favour the novel option that is usually a good, non-specific indicator of new information. This inability might result from a tendency towards distributed attention (although distributed attention might as well be the outcome of this difficulty), but also might reflect memory constraints in remembering the calculated values and comparing them. Indeed, our findings show that lower working memory capacity results in larger novelty preference, while also clearly showing an effect of cognitive constraints leading children to make more choices based on novelty. However, an inability to remember the values could result in choosing the novelty option randomly, as a "mistake' when they have to choose quickly and cannot incorporate the values of the other options to their action plans, and not necessarily because of its appeal. Although we cannot exclude this possibility, especially since we did not include a "random", control option, lower inhibition is also involved in more novelty choices, suggesting not only that children could not keep the other options' values in mind as a priority, but that they indeed found the novel options attractive –

and tried, even if not successfully, to suppress their distraction. Finally, higher novelty valuation could also result from an overestimation of volatility in younger ages – as novelty is considered optimal for learning in changing environments, although we do not examine this possibility in our studies.

5.2.2. Interaction between goal-following and reward values modulate adult decision-making

Adult participants in our experiments in Chapter 3 have distinct behaviours as compared to the younger groups; they seem to prefer both external and informational rewards equally. This is the case even when they are under time pressure, suggesting that they consider both options equally valuable. However, our task design in the first two experiments of Chapter 3 involved following specific steps to acquire either an external reward or fully resolve perceptual uncertainty, even without explicit instructions to do so (i.e., it presupposed goal-directed action and ability to follow steps), while rewarding feelings were only experienced in each trial because the goal was approaching (in the external reward case) or due to partial perceptual relief of uncertainty. This was not the case in our final experiment, where both types of rewards were immediately available and not the result of goal-following. The fact that adults were no longer attracted by these options more than the novelty one is suggestive that the pursuit of a goal was acting as a stronger motivation to them, adding value to these options. While in the case of the external reward adults were possibly not motivated by the specific secondary reward-by-association, the fact that perceptual uncertainty was also not that appealing to them without a goal is more interesting. This points towards a direction that goals themselves hold intrinsic value, often initially associated with other specific (external or internal) rewards but then playing an independent role and they can motivate humans similarly. Such a point was recently made by Molinaro and Collins (2023), who suggest that rewards should be viewed as conditional to goals – even primary rewards such as food are not rewarding when a person is not hungry. In our results, higher order goals were probably important to adults, such as "winning" as a general value, and completing a puzzle – for the sake of completion.

5.2.3. Complexity has different effects on preschoolers’ exploration and explicit preferences

As it is evident from Chapters 2 and 4, increasing an object’s visual and haptic complexity affects how much time is spent on its visual processing and its manipulation, specifically when complexity is manipulated through the number and organisation of the object’s elements. This effect seems to be linear in the visual domain, but not clear in the haptic one – possibly affected by factors such as haptic efficiency and experience. The documented linear relationship in the visual domain extends previous findings in the adult and developmental literature, by showing that children can track the amount of information held by objects and adjust their attention to accurately process them. The fact that we found a linear and not an inverted U-shaped curve relationship, as previously suggested in developmental studies, probably has to do with the learnability of the specific object characteristics for the observers – the high complexity objects were still not too complex for preschoolers to process. Furthermore, no previous study in this age group has explicitly manipulated visual complexity, instead of predictability or incongruity in the temporal or conceptual domain. The fact that the haptic exploration time relationship to complexity approximated an inverted U- shaped curve, further suggests that preschoolers explore longer only if they can process the available information.

In contrast, complexity did not consistently influence explicit preference – participants had large individual differences in preferences and were also not always consistent across modalities. This is consistent with recent studies emphasising the need for analysing participants’ individual predispositions when it comes to complexity preferences. Furthermore, our findings should be considered from a methodological perspective too, showing that exploration time and explicit preference should be used as separate measures for different processes in participants.

5.3. Methodological contributions

5.3.1. Mouse-tracking tasks for value-based decisions in children

Mouse and finger-tracking methodology is becoming increasingly used in psychological experiments, providing deeper insights into cognitive processes and enabling measures to move away from just accuracy or reaction times. Cognitive psychological paradigms are largely still focused on language and reasoning research, while value-based paradigms are mostly thriving in marketing and consumer choice research. Moreover, both fields are mostly creating adult-focused tasks. To our knowledge, value-based mouse-tracking tasks for children only exist in dietary research, studying children older than 7 years old (Pearce et al., 2020).

Our work provides an example of a child-friendly decision-making task, applying the classic 2AFC format to value-based decisions. Comparison of kinematics results from the mouse-tracking and finger-tracking procedures (the latter of which we utilised with younger kids) showed high consistency and thus paves the road to study decisions in younger kids who have very noisy mouse data, or cannot use the mouse. Importantly, when interested in goal-directed action, value comparison and conflict, movement tracking can prove much more suitable than eye-tracking (which is the most common process-tracing methodology used with children) as it focuses on action plans, while eye movements are largely driven bottom-up by available information in the environment – although movement trajectories are not immune to such effects either.

Regarding the kinematics data, our study showed that commonly used geometrical and entropy features can convey information about conflict and about preferences and valuation of options. However, since the theoretical assumptions about kinematic features are mostly based on options being compared on the basis of right/wrong answer, the most accurate features are still to be determined in value-based decisions. While we utilised the most common measures to allow for comparisons, analyses of temporal measures or specific hypotheses about stages of the decision process should be explored too.

5.3.2. Quantitative manipulation of object complexity

The difficulty in manipulating visual complexity has resulted in great inconsistency in studying both psychological and aesthetic aspects of human response to this property. Studies have very often created stimuli manipulating the amount and organisation of elements in uncontrolled ways (as did we in our Chapter 2 experiment), or their symmetry, or used richer stimuli such as photos and paintings whose complexity has proved even harder to measure consistently. More recent studies have attempted to use computational methods to directly manipulate visual complexity or classify stimuli based on it. Examples of such methods include the compressibility of images, the algorithmic manipulation of the elements' combinations, such as patterns of triangles and squares, or the extraction of the skeleton of shapes using computational geometry. We adopted a similar approach to our study, using the amount of information of visual stimuli, as quantified in their GIF compressibility. While not necessarily a better approach than the other computational ones mentioned, such an objective method can provide more reliable results and possibly be more comparable to other studies and can generate large samples of test stimuli. Indeed, in our case, the visual exploration of objects by children seemed consistent with previous findings, suggesting that our complexity manipulation was efficient – although even more complex or asymmetric shapes might have afforded more interesting observations. However, when haptic complexity also has to be quantified simultaneously in the same objects, using visual complexity measures might not be optimal because object manipulation makes different information available to the hands. As a result, methods such as the internal skeleton, applied to 3D shapes, might prove better candidates.

5.4. Limitations

We discussed the specific methodological limitations to our studies in the separate experimental chapters. We therefore discuss limitations here from a broader perspective. An important limitation regards the manipulation of independent variables in our series of mouse-tracking experiments, which limits the conclusions we can draw from comparing the experiments' results. This issue primarily regards the operationalisation of the concepts of external reward and perceptual uncertainty, which were confounded with goal-following in our first three experiments. Even the concept of novelty, which was better operationalised, still allowed for some level of category prediction (e.g., expectations for cat cartoons, albeit a new one every time) in the first two experiments. Thus novelty was probably slightly confounded with expectations for uncertainty resolution. The fact that these concepts were de-confounded in the final two experiments and still lead to similar observations suggests that they adequately captured differences – or that they lead to interesting findings such as the effect of goals for adults. However, the simultaneous manipulation of other factors such as the decision time-limit, the decision horizon, the sample age range (e.g., Experiment 3 only included children) makes final conclusions partial and in need for more specific separate experiments to clarify each factor's effect. Similar issues arise from the inclusion of different EF tasks in the different experiments, which yielded different results. The lack of a working memory task in our final experiment also prevented us from drawing more comprehensive conclusions about the role of EF in exploratory choices.

Furthermore, although our variables captured differences between uncertainty- and novelty-based exploration, the concepts themselves were not quantified. The fact that they were unrelated to a central external goal allowed us ot explore the possible differences between different simultaneous motives of closing specific information gaps and sensation-seeking, which are naturalistic and occur in everyday life decisions, but cannot directly inform research on the explore-exploit dilemma -- especially regarding the separate contribution of uncertainty and novelty as informational attributes in a specific problem-solving context. Our experiments could possibly make the case that children's directed exploration is not that consistent and they are also switching more, but a fourth option including no motive (i.e., not a novelty one) would be needed to dissociate between novelty and random choices.

Moreover, our mouse-tracking task data showed great variability in terms of individual participants' choices in all age groups – but particularly in children and adults. Since we only present results at the aggregate level, it is possible that we are missing explanatory factors that could affect preferences other than EF proficiency.

Furthermore, our object complexity experiments' main limitation was the manipulation of complexity in the haptic domain. This was better uncontrolled in Chapter 2, but in Chapter 4 the haptic complexity manipulation was based on measured visual characteristics (although piloting allowed for some control on the experienced level of complexity). There was no control of other factors that might affect haptic expertise, such as multisensory development and working memory (in our final experiment).

5.5. Future directions

5.5.1. Novelty as sensation-seeking vs. novelty as information

Our studies suggested that children are more novelty-driven than adolescents and adults in the sense that, apart from aiming to cover information gaps (i.e., keeping track of missing information and directing their exploration based on learning), they also seek sensory stimulation aiming to derive pleasure from broader learning goals or generally to generate arousal. However, novelty is also often conceptualised as an informational attribute, capturing whether a stimulus has been explored before or how much time has passed since an option has been explored (this refers to environments where volatility is known or hypothesised). It becomes apparent that the two conceptualisations of novelty-seeking are dependent on its instrumentality – if all options in the environment are possible candidates to convey information for the acquisition of a reward, novelty is a useful attribute to keep track of and follow, and choosing it is subject to previously studied factors, such as the observer's or the environment's state of knowledge and available information. If some options in this environment are novel but not instrumental to the goal, choosing them more likely reflects a drive towards the pleasure of learning, broader curiosity or boredom from the main quest, and possibly relates to sensation-seeking in a more general sense (but this is another possible hypothesis to be tested). An attempt to dissociate between these different motives for information-seeking would be interesting and can be studied in a task where all the aforementioned variables would be present: a main goal, options with different informativity (learnable sequences and novel options of the same learnable category) as well as novel options with no function apart from new sensory stimuli. A gamified task such as the ones used by Nussenbaum et al. (2022) or Poli et al. (2022) would allow for the simultaneous presentation of different options (e.g., as different characters presenting information with different probabilities and new characters with unknown probabilities, as well as new characters presenting no relevant information at all). Different predictions would allow for exploration based on different factors. Furthermore, apart from quantifying uncertainty-based exploration, novelty-based exploration and sensation-seeking, such a task would allow us to control for random exploration, as unpredictable switches between the other options – and therefore dissociate between such behaviour and looking for new learning goals. Developing a task suitable across different age groups would also allow us to identify both developmental and individual differences in exploratory motives.

5.5.2. Neural mechanisms of uncertainty-based and novelty-based exploration across development

Following on from the aforementioned idea, different exploratory motives and behaviours have been proposed to be supported by different systems in the brain (the dopaminergic system proposedly supports informational reward seeking, while the opioid system is responsible for pleasure; see the wanting vs. liking information distinction made by Litman, 2005). Furthermore, it has been suggested that exploration based on uncertainty is also supported by different brain areas and by different neuromodulators, depending on the amount of available information in each context. To our knowledge, the neural substrates of these different exploratory behaviours have only been studied in adults, consistently showing distinct activations of areas of the prefrontal cortex. For example, two recent studies using fNIRS and fMRI (Li et al., 2019; Tomov et al., 2020) have shown that relative uncertainty leads to activation in the right rostrolateral prefrontal cortex and drives directed exploration, while total uncertainty affects the right dorsolateral prefrontal cortex and drives random exploration. Different neuromodulators have also been shown to influence random vs. uncertainty-based exploration (Dubois, Habicht, Michely, Moran, Dolan, & Hauser, 2021). In the relative vs. total uncertainty studies, novelty and uncertainty are used interchangeably, as both regard instrumental contexts and reflect the informativity of a stimulus. A task usable across different age groups which would further dissociate between these variables and non-instrumental options would possibly shed light in resolving uncertainty vs. pleasure motives. However, since both fNIRS and fMRI task designs have specific methodological requirements, a visually rich gamified task with free choice would possibly be unsuitable; a 2AFC tasks comparing between different options would be preferable. To also measure the effects of different neuromodulators – especially noradrenaline – we could combine pupil dilation tracking to our methodology.

5.5.3. Effects of cognitive control on uncertainty and novelty preference

Our results from the Chapter 3 show an effect of cognitive control on participants' exploratory preferences. As a further extension of the suggested tasks, it would be interesting to include age-specific EF tasks. The more accurate measurement of exploratory preferences associated with cognitive control measures would clarify whether the effects in our experiments were due to participants with lower control being more prone to novelty/sensation-seeking or they could not keep track of option values due to memory/general cognitive constraints.

5.5.4. Relationship between imagined object manipulation and motor activation

Both Chapter 2 and Chapter 4 touch on the topic of spatial and motor imagery and how it is affected by object complexity. However, they only indirectly measure it (through fitting accuracy) or presuppose its involvement in the tasks. In the preschool years, imagined object manipulation is probably greatly affected by immature multisensory integration, such that an underdeveloped ability to combine visual, haptic and proprioceptive information into coherent schemas might render the mental manipulation of objects harder or wrong. This might be even more pronounced as object complexity increases, and difficulty to mentally represent an object might lead children to activate motor plans to support their imagery. This was partially our initial hypothesis in Chapter 2. However, it might also be possible that motor activations are subtler than complete representational gestures, so measuring hand motor activation through EMG might have been more powerful at revealing effects. Incorporating such a methodology would allow us to observe if children spontaneously support their imagined manipulation with their hands (as an embodied cognition perspective would suggest). Furthermore, a measurement of their ability to integrate information from the senses through a suitable task (e.g., a reach-to-grasp task accuracy) would show whether the effect of motor activation is stronger in less accurate children. Since multisensory development spans early and middle childhood, older participants should also be tested in such an experiment.

5.6. Conclusion

This thesis proposed new experimental approaches in studying children's, adolescents' and adults' exploratory behaviour and relationship with information across development.

Using computerised tasks, we observed that some children show an increased preference for perceptual novelty. We suggest that this tendency relates to their cognitive control maturation, either as a compromised ability to keep track of more specific informative goals, or as a difficulty to inhibit sensation-seeking tendencies. We make a more theoretical case for novelty often being treated both as an informational attribute to achieve a goal and as an attribute to increase arousal, obscuring its effects on decision-making. Moreover, using physical objects, we show that young children and preschoolers keep track of visual and haptic object complexity and adjust their exploration and manipulation time based on it. We also differentiate between exploration time as a measure of informational processing and explicit preferences for complex objects. Our thesis also tried to contribute on the methodological level, by designing and applying kinematic analyses on child-friendly tasks.

We hope our theoretical, experimental and methodological contributions will help future research on exploration across development.

Appendices

Appendix A – Chapter 3 complementary tables

A.1. Pairwise comparisons, Option*Experiment, full sample

Option	(I) Experiment	(J) Experiment	Mean Difference (I-J)	Std. Error	Sig.[d]	95% Confidence Interval for Difference[d]	
						Lower Bound	Upper Bound
ER	1.00	2.00	-2.015*	.618	.008	-3.663	-.367
		3.00	-1.331[b]	.916	.887	-3.773	1.111
		4.00	-1.618	.680	.110	-3.431	.195
	2.00	1.00	2.015*	.618	.008	.367	3.663
		3.00	.684[b]	.920	1.000	-1.768	3.137
		4.00	.397	.685	1.000	-1.430	2.225
	3.00	1.00	1.331[c]	.916	.887	-1.111	3.773
		2.00	-.684[c]	.920	1.000	-3.137	1.768
		4.00	-.287[c]	.963	1.000	-2.854	2.279
	4.00	1.00	1.618	.680	.110	-.195	3.431
		2.00	-.397	.685	1.000	-2.225	1.430
		3.00	.287[b]	.963	1.000	-2.279	2.854
IR	1.00	2.00	-3.637*	.558	<.001	-5.123	-2.150
		3.00	-2.317*,[b]	.826	.033	-4.520	-.114
		4.00	-2.915*	.613	<.001	-4.551	-1.280
	2.00	1.00	3.637*	.558	<.001	2.150	5.123
		3.00	1.320[b]	.830	.681	-.893	3.533
		4.00	.721	.618	1.000	-.928	2.370
	3.00	1.00	2.317*,[c]	.826	.033	.114	4.520
		2.00	-1.320[c]	.830	.681	-3.533	.893
		4.00	-.598[c]	.869	1.000	-2.914	1.717
	4.00	1.00	2.915*	.613	<.001	1.280	4.551
		2.00	-.721	.618	1.000	-2.370	.928
		3.00	.598[b]	.869	1.000	-1.717	2.914

NS	1.00	2.00	-2.882*	.587	<.001	-4.446	-1.317
		3.00	-2.453*,b	.870	.032	-4.772	-.135
		4.00	-3.657*	.646	<.001	-5.378	-1.936
	2.00	1.00	2.882*	.587	<.001	1.317	4.446
		3.00	.429b	.873	1.000	-1.900	2.758
		4.00	-.775	.651	1.000	-2.510	.960
	3.00	1.00	2.453*,c	.870	.032	.135	4.772
		2.00	-.429c	.873	1.000	-2.758	1.900
		4.00	-1.204c	.914	1.000	-3.641	1.233
	4.00	1.00	3.657*	.646	<.001	1.936	5.378
		2.00	.775	.651	1.000	-.960	2.510
		3.00	1.204b	.914	1.000	-1.233	3.641
Based on estimated marginal means							
*. The mean difference is significant at the .05 level.							
b. An estimate of the modified population marginal mean (J).							
c. An estimate of the modified population marginal mean (I).							
d. Adjustment for multiple comparisons: Bonferroni.							

A.2. Pairwise comparisons, Age Group* Option* Experiment

Age group	Option	(I) Experiment	(J) Experiment	Mean Difference (I-J)	Std. Error	Sig.d	95% Confidence Interval for Differenced	
							Lower Bound	Upper Bound
Children	ER	1.00	2.00	-1.441	1.049	1.000	-4.238	1.357
			3.00	-1.317	1.091	1.000	-4.225	1.591
			4.00	-1.875	1.076	.498	-4.743	.993
		2.00	1.00	1.441	1.049	1.000	-1.357	4.238
			3.00	.124	1.102	1.000	-2.814	3.062

			4.00	-.434	1.087	1.000	-3.332	2.464
		3.00	1.00	1.317	1.091	1.000	-1.591	4.225
			2.00	-.124	1.102	1.000	-3.062	2.814
			4.00	-.558	1.127	1.000	-3.563	2.447
		4.00	1.00	1.875	1.076	.498	-.993	4.743
			2.00	.434	1.087	1.000	-2.464	3.332
			3.00	.558	1.127	1.000	-2.447	3.563
	IR	1.00	2.00	-3.095*	.947	.008	-5.619	-.571
			3.00	-2.690*	.984	.041	-5.314	-.066
			4.00	-2.593*	.971	.049	-5.181	-.005
		2.00	1.00	3.095*	.947	.008	.571	5.619
			3.00	.405	.994	1.000	-2.246	3.056
			4.00	.502	.981	1.000	-2.113	3.117
		3.00	1.00	2.690*	.984	.041	.066	5.314
			2.00	-.405	.994	1.000	-3.056	2.246
			4.00	.097	1.017	1.000	-2.614	2.809
		4.00	1.00	2.593*	.971	.049	.005	5.181
			2.00	-.502	.981	1.000	-3.117	2.113
			3.00	-.097	1.017	1.000	-2.809	2.614
	NS	1.00	2.00	-5.032*	.996	<.001	-7.688	-2.376
			3.00	-2.715	1.036	.057	-5.476	.047

			4.00	-2.093	1.021	.251	-4.817	.630
		2.00	1.00	5.032*	.996	<.001	2.376	7.688
			3.00	2.318	1.046	.168	-.472	5.107
			4.00	2.939*	1.032	.029	.187	5.691
		3.00	1.00	2.715	1.036	.057	-.047	5.476
			2.00	-2.318	1.046	.168	-5.107	.472
			4.00	.621	1.070	1.000	-2.232	3.475
		4.00	1.00	2.093	1.021	.251	-.630	4.817
			2.00	-2.939*	1.032	.029	-5.691	-.187
			3.00	-.621	1.070	1.000	-3.475	2.232
Adolescents	ER	1.00	2.00	-1.500	1.142	.571	-4.257	1.257
			4.00	-1.908	1.375	.500	-5.228	1.412
		2.00	1.00	1.500	1.142	.571	-1.257	4.257
			4.00	-.408	1.375	1.000	-3.728	2.912
		4.00	1.00	1.908	1.375	.500	-1.412	5.228
			2.00	.408	1.375	1.000	-2.912	3.728
	IR	1.00	2.00	-4.548*	1.030	<.001	-7.036	-2.061
			4.00	-3.294*	1.240	.026	-6.290	-.299
		2.00	1.00	4.548*	1.030	<.001	2.061	7.036
			4.00	1.254	1.240	.940	-1.742	4.250
		4.00	1.00	3.294*	1.240	.026	.299	6.290

			2.00	-1.254	1.240	.940	-4.250	1.742
	NS	1.00	2.00	-3.456*	1.084	.005	-6.074	-.838
			4.00	-4.657*	1.305	.001	-7.810	-1.505
		2.00	1.00	3.456*	1.084	.005	.838	6.074
			4.00	-1.201	1.305	1.000	-4.354	1.951
		4.00	1.00	4.657*	1.305	.001	1.505	7.810
			2.00	1.201	1.305	1.000	-1.951	4.354
Adults	ER	1.00	2.00	-3.105*	1.017	.008	-5.560	-.650
			4.00	-1.071	1.056	.935	-3.620	1.479
		2.00	1.00	3.105*	1.017	.008	.650	5.560
			4.00	2.034	1.076	.180	-.564	4.633
		4.00	1.00	1.071	1.056	.935	-1.479	3.620
			2.00	-2.034	1.076	.180	-4.633	.564
	IR	1.00	2.00	-3.266*	.917	.001	-5.482	-1.051
			4.00	-2.859*	.952	.009	-5.159	-.559
		2.00	1.00	3.266*	.917	.001	1.051	5.482
			4.00	.407	.971	1.000	-1.937	2.751
		4.00	1.00	2.859*	.952	.009	.559	5.159
			2.00	-.407	.971	1.000	-2.751	1.937

	NS	1.00	2.00	-.157	.965	1.000	-2.488	2.174
			4.00	-4.220*	1.002	<.001	-6.641	-1.799
		2.00	1.00	.157	.965	1.000	-2.174	2.488
			4.00	-4.063*	1.021	<.001	-6.530	-1.596
		4.00	1.00	4.220*	1.002	<.001	1.799	6.641
			2.00	4.063*	1.021	<.001	1.596	6.530

Based on estimated marginal means

*. The mean difference is significant at the .05 level.

a. The level combination of factors in (J) is not observed.

b. The level combination of factors in (I) is not observed.

d. Adjustment for multiple comparisons: Bonferroni.

Appendix B – Chapter 4 complementary results

B.1. Mixed linear models predicting ET from participants' clusters

Participants' cluster in Visual Liking did not predict their visual ET of the low-complexity object ($\chi2(2) = 1.081$, $p = .582$), the medium complexity object ($\chi2(2) = 0.742$, $p = .690$) or the high complexity object ($\chi2(2) = 2.462$, $p = .292$). Similarly, participants' cluster in Visual Play did not predict their visual ET of the low-complexity object ($\chi2(5) = 3.114$, $p = .682$), the medium complexity object ($\chi2(5) = 3.252$, $p = .661$) or the high complexity object ($\chi2(5) = 1.403$, $p = .924$).
Participants' cluster in Haptic Liking did not predict their haptic ET of the low-complexity object ($\chi2(5) = 0.710$, $p = .983$), the medium complexity object ($\chi2(5) = 1.956$, $p = .855$) or the high complexity object ($\chi2(5) = 0.128$, $p = .999$). In a similar pattern, participants' cluster in Haptic Play did not predict their haptic ET of the low-complexity object ($\chi2(5) = 1.442$, $p = .950$), the medium complexity object ($\chi2(5) = 2.451$, $p = .784$) or the high complexity object ($\chi2(5) = 3.311$, $p = .652$).

References

Addyman, C., & Mareschal, D. (2013). Local redundancy governs infants' spontaneous orienting to visual-temporal sequences. *Child development*, *84*(4), 1137-1144.

Addyman, C., Rocha, S., Fautrelle, L., French, R. M., Thomas, E., & Mareschal, D. (2017). Embodiment and the origin of interval timing: kinematic and electromyographic data. *Experimental brain research*, *235*, 923-930.

Adolph, K. E. (2000). Specificity of learning: Why infants fall over a veritable cliff. *Psychological Science,* 11, 290-295.

Aguirre, M., Brun, M., Couderc, A., Reboul, A., Senez, P., & Mascaro, O. (2022). Knowledge in Sight: Toddlers Plan Efficient Epistemic Actions by Anticipating Learning Gains. *Cognitive Science*, *46*(2).

Aitken, P. P. (1974). Judgments of pleasingness and interestingness as functions of visual complexity. *Journal of Experimental Psychology*, *103*(2), 240.

Alberti, E. T., & Witryol, S. L. (1990). Children's preference for complexity as a function of perceived units in collative motivation. *The Journal of genetic psychology*, *151*(1), 91-101.

Amedi, A., Jacobson, G., Hendler, T., Malach, R., & Zohary, E. (2002). Convergence of visual and tactile shape processing in the human lateral occipital complex. *Cerebral cortex*, *12*(11), 1202-1212.

Ariely, D., & Zakay, D. (2001). A timely account of the role of duration in decision making. *Acta psychologica*, *108*(2), 187-207.

Attneave, F. (1954). Some informational aspects of visual perception. *Psychological review*, *61*(3), 183.

Attneave, F. (1957). Physical determinants of the judged complexity of shapes. *Journal of experimental Psychology*, *53*(4), 221.

Badre, D., Doll, B. B., Long, N. M., & Frank, M. J. (2012). Rostrolateral prefrontal cortex and individual differences in uncertainty-driven exploration. *Neuron*, *73*(3), 595-607.

Baldassarre, G. (2011, August). What are intrinsic motivations? A biological perspective. In *2011 IEEE international conference on development and learning (ICDL)*(Vol. 2, pp. 1-8). IEEE.

Baranes, A. F., Oudeyer, P. Y., & Gottlieb, J. (2014). The effects of task difficulty, novelty and the size of the search space on intrinsically motivated exploration. *Frontiers in neuroscience*, *8*, 317.

Barrett, T. M., & Needham, A. (2008). Developmental differences in infants' use of an object's shape to grasp it securely. *Developmental Psychobiology*, 50, 97-106.

Barsalou, L. W. (2016). On staying grounded and avoiding quixotic dead ends. *Psychonomic bulletin & review*, *23*(4), 1122-1142.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting linear mixed-effects models using lme4. *arXiv preprint arXiv:1406.5823*.

Benedetti, V., Gronchi, G., Gavazzi, G., Bravi, R., Grasso, S., Giovannelli, F., & Viggiano, M. P. (2021). Awareness motor intention and inhibitory control: the role of reactive and proactive components. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 43, No. 43).

Berlyne, D. E. (1954). A theory of human curiosity. British Journal of Psychology. 1954;45:180-91

Berlyne, D. E. (1960). Conflict, arousal, and curiosity. London: McGraw-Hill; 1960.

Berlyne, D. E., & Lawrence, G. H. (1964). Effects of complexity and incongruity variables on GSR, investigatory behavior, and verbally expressed preference. *The Journal of General Psychology*, *71*(1), 21-45.

Berlyne, D. E. (1966). Curiosity and Exploration: Animals spend much of their time seeking stimuli whose significance raises problems for psychology. *Science*, *153*(3731), 25-33.

Berlyne, D. E. (1970). Novelty, complexity, and hedonic value. *Perception & psychophysics*, *8*(5), 279-286.

Berlyne, D. E., & Crozier, J. B. (1971). Effects of complexity and prechoice stimulation on exploratory choice. *Perception & Psychophysics*, *10*(4), 242-246.

Berlyne, D. E., Ogilvie, J. C., & Parham, L. C. (1968). The dimensionality of visual complexity, interestingness, and pleasingness. *Canadian Journal of Psychology / Revue canadienne de psychologie, 22*(5), 376–387.

Berridge, K. C. (1999). Pleasure, pain, desire, and dread: Hidden core processes of emotion. In D. Kahneman, E. Diener, & N. Schwarz, (Eds.), *Well-being: Foundations of hedonic psychology* (pp. 527±559). New York: Russell Sage Foundation.

Berridge, K. C., & Robinson, T.E. (1998). The role of dopamine in reward: Hedonics, learning, or incentive salience? Brain Research Reviews, 28, 308±367.

Berthier, N. E., & Carrico, R. L. (2010). Visual information and object size in infant reaching. *Infant Behavior and Development*, 33, 555-566.

Betsch, T., Haberstroh, S., Molter, B., & Glöckner, A. (2004). Oops, I did it again—relapse errors in routinized decision making. *Organizational behavior and human decision processes*, *93*(1), 62-74.

Bezdjian, S., Baker, L. A., Lozano, D. I., & Raine, A. (2009). Assessing inattention and impulsivity in children during the Go/NoGo task. *British Journal of Developmental Psychology*, *27*(2), 365-383.

Blain, B., & Sharot, T. (2021). Intrinsic reward: potential cognitive and neural mechanisms. *Current Opinion in Behavioral Sciences*, *39*, 113-118.

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.

Blanco, N. J., & Sloutsky, V. M. (2021). Systematic exploration and uncertainty dominate young children's choices. *Developmental Science*, *24*(2), e13026.

Blanco, N. J., Turner, B. M., & Sloutsky, V. M. (2023). The benefits of immature cognitive control: How distributed attention guards against learning traps. *Journal of Experimental Child Psychology*, *226*, 105548. https://doi.org/10.1016/j.jecp.2022.105548

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

Boncoddo, R., Dixon, J. A., & Kelley, E. (2010). The emergence of a novel representation from action: evidence from preschoolers. *Developmental Science*, *13*(2), 370-377.

Brändle, F., Wu, C. M., & Schulz, E. (2020). What Are We Curious about? *Trends in Cognitive Sciences*, *24*(9), 685–687. https://doi.org/10.1016/j.tics.2020.05.010

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.

Bromberg-Martin, E. S., & Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron*, *63*(1), 119-126.

Chamorro-Premuzic, T., Burke, C., Hsu, A., & Swami, V. (2010). Personality predictors of artistic preferences as a function of the emotional valence and perceived complexity of paintings. *Psychology of Aesthetics, Creativity, and the Arts*, *4*(4), 196.

Cappuccio, M. L., Chu, M., & Kita, S. (2013). Pointing as an instrumental gesture: Gaze representation through indication. *Humana. Mente: Journal of Philosophical Studies*.

Carteron, A. M. (2022) *Development of routine and supervisory processes in sequential action control*. PhD thesis, Birkbeck, University of London.

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.

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

Chen, X., Twomey, K., & Westermann, G. (2022). The role of metacognitive abilities and curiosity in learning.*PsyArxiv*.
https://psyarxiv.com/asqkj/download?format=pdf

Chipman, S. F. (1977). Complexity and structure in visual patterns. *Journal of Experimental Psychology: General*, *106*(3), 269.

Chipman, S. F. (1977). Complexity and structure in visual patterns. *Journal of Experimental Psychology: General*, *106*(3), 269.

Chu, M., & Kita, S. (2011). The nature of gestures' beneficial role in spatial problem solving. *Journal of Experimental Psychology: General*, *140*(1), 102.

Chu, M., & Kita, S. (2016). Co-thought and co-speech gestures are generated by the same action generation process. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42, 257–270.

Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annual review of neuroscience*, *33*, 269-298.

Cogliati Dezza, I., Yu, A. J., Cleeremans, A., & Alexander, W. (2017). Learning the value of information and reward over time when solving exploration-exploitation problems. *Scientific reports*, *7*(1), 16919.

Cohen, L. B., DeLoache, J. S., & Rissman, M. W. (1975). The effect of stimulus complexity on infant visual attention and habituation. *Child Development*, 611-617.

Cockburn, J., Man, V., Cunningham, W., & O'Doherty, J. P. (2021). *Novelty and uncertainty interact to regulate the balance between exploration and exploitation in the human brain* [Preprint]. Neuroscience. https://doi.org/10.1101/2021.10.13.464279

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1481), 933-942.

Colas, C., Fournier, P., Chetouani, M., Sigaud, O., & Oudeyer, P. Y. (2019, May). Curious: intrinsically motivated modular multi-goal reinforcement learning. In *International conference on machine learning* (pp. 1331-1340). PMLR.

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

Coughlin, C., Hembacher, E., Lyons, K. E., & Ghetti, S. (2015). Introspection on uncertainty and judicious help-seeking during the preschool years. *Developmental Science*, *18*(6), 957-971.

Crupi, V., Nelson, J. D., Meder, B., Cevolani, G., & Tentori, K. (2018). Generalized information theory meets human cognition: Introducing a unified framework to model uncertainty and information search. *Cognitive Science*, *42*(5), 1410-1456.

Dale, R., & Duran, N. D. (2011). The cognitive dynamics of negated sentence verification. *Cognitive science*, *35*(5), 983-996.

Day, H. Y. (1967). Evaluations of subjective complexity, pleasingness and interestingness for a series of random polygons varying in complexity. *Perception & Psychophysics*, *2*, 281-286.

Day, H. (1968). The importance of symmetry and complexity in the evaluation of complexity, interest and pleasingness. *Psychonomic Science*, *10*(10), 339-340.

Day , H. I. (1971). " The measurement of specific curiosity " . In *Intrinsic Motivation: A New Direction in Education* , Edited by: Day , H. I. , Berlyne , D. E. and Hunt , D. E. 99 – 112 . Toronto : Holt, Rinehart & Winston .

de Mooij, S. M., Raijmakers, M. E., Dumontheil, I., Kirkham, N. Z., & van der Maas, H. L. (2021). Error detection through mouse movement in an online adaptive learning environment. *Journal of Computer Assisted Learning*, *37*(1), 242-252.

Delgado, B., Gómez, J. C., & Sarriá, E. (2011). Pointing gestures as a cognitive tool in young children: experimental evidence. *Journal of Experimental Child Psychology*, 110, 299–312.

Dember, W. N., & Earl, R. W. (1957). Analysis of exploratory, manipulatory, and curiosity behaviors. *Psychological review*, *64*(2), 91.

Desender, K., Boldt, A., & Yeung, N. (2018). Subjective confidence predicts information seeking in decision making. *Psychological science*, *29*(5), 761-778.

Desender, K., Murphy, P., Boldt, A., Verguts, T., & Yeung, N. (2019). A postdecisional neural marker of confidence predicts information-seeking in decision-making. *Journal of Neuroscience*, *39*(17), 3309-3319.

Dezza, I. C., Maher, C., & Sharot, T. (2022). People adaptively use information to improve their internal states and external outcomes. *Cognition*, *228*, 105224.

Di Domenico, S. I., & Ryan, R. M. (2017). The emerging neuroscience of intrinsic motivation: A new frontier in self-determination research. *Frontiers in human neuroscience*, *11*, 145.

Dieciuc, M. A., Roque, N. A., & Boot, W. R. (2019). The spatial dynamics of mouse-tracking reveal that attention capture is stimulus-driven rather than contingent upon top-down goals. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(10), 1285.

Dotan, D., & Dehaene, S. (2013). How do we convert a number into a finger trajectory?. *Cognition*, *129*(3), 512-529.

Dotan, D., Meyniel, F., & Dehaene, S. (2018). On-line confidence monitoring during decision making. *Cognition*, *171*, 112-121.

Drechsler, R., Rizzo, P., & Steinhausen, H. C. (2010). Decision making with uncertain reinforcement in children with attention deficit/hyperactivity disorder (ADHD). *Child Neuropsychology*, *16*(2), 145-161.

Dubey, R., & Griffiths, T. L. (2020). Understanding exploration in humans and machines by formalizing the function of curiosity. *Current Opinion in Behavioral Sciences*, *35*, 118–124. https://doi.org/10.1016/j.cobeha.2020.07.008

Duran, N. D., Dale, R., Kello, C. T., Street, C. N., & Richardson, D. C. (2013). Exploring the movement dynamics of deception. *Frontiers in psychology*, *4*, 140.

Ekman, P. (1999). Basic emotions. *Handbook of cognition and emotion*, *98*(45-60), 16.

Eisenman, R., & Gellens, H. K. (1968). Preferences for complexity-simplicity and symmetry-asymmetry. *Perceptual and Motor Skills*, *26*(3), 888-890.

Erb, C. D., Moher, J., Sobel, D. M., & Song, J. H. (2016). Reach tracking reveals dissociable processes underlying cognitive control. *Cognition*, *152*, 114-126.

Erb, C. D., Moher, J., Song, J. H., & Sobel, D. M. (2018). Reach tracking reveals dissociable processes underlying inhibitory control in 5-to 10-year-olds and adults. *Developmental science*, *21*(2), e12523.

Erb, C. D., Moher, J., Song, J. H., & Sobel, D. M. (2017). Cognitive control in action: Tracking the dynamics of rule switching in 5-to 8-year-olds and adults. *Cognition*, *164*, 163-173.

Faulkenberry, T. J., Montgomery, S. A., & Tennes, S. A. N. (2015). Response trajectories reveal the temporal dynamics of fraction representations. *Acta Psychologica*, *159*, 100-107.

Fantz, R. L. (1964). Visual experience in infants: Decreased attention to familiar patterns relative to novel ones. *Science*, *146*(3644), 668-670.

Farmer, T. A., Cargill, S. A., Hindy, N. C., Dale, R., & Spivey, M. J. (2007). Tracking the continuity of language comprehension: Computer mouse trajectories suggest parallel syntactic processing. *Cognitive science*, *31*(5), 889-909.

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. https://doi.org/10.1016/j.cobeha.2020.05.014

Forestier, S., Mollard, Y., Caselli, D., & Oudeyer, P. Y. (2016, December). Autonomous exploration, active learning and human guidance with open-source

Poppy humanoid robot platform and Explauto library. In *The Thirtieth Annual Conference on Neural Information Processing Systems (NIPS 2016).*

Fragaszy, D. M., Kuroshima, H., & Stone, B. W. (2015). “Vision for action” in young children aligning multi-featured objects: Development and comparison with nonhuman primates. *PLoS One*, 10.

Freud, S. (1915). Instincts and their vicissitudes. *Standard edition*, *14*(1957), 109-40.

Freeman, J. B., & Ambady, N. (2009). Motions of the hand expose the partial and parallel activation of stereotypes. *Psychological science*, *20*(10), 1183-1188.

Freeman, J. B., & Ambady, N. (2010). MouseTracker: Software for studying real-time mental processing using a computer mouse-tracking method. *Behavior research methods*, *42*(1), 226-241.

Freeman, J. B., Ma, Y., Han, S., & Ambady, N. (2013). Influences of culture and visual context on real-time social categorization. *Journal of experimental social psychology*, *49*(2), 206-210.

Freeman, J. B., Pauker, K., & Sanchez, D. T. (2016). A perceptual pathway to bias: Interracial exposure reduces abrupt shifts in real-time race perception that predict mixed-race bias. *Psychological Science*, *27*(4), 502-517.

Frick, A., Hansen, M. A., & Newcombe, N. S. (2013). Development of mental rotation in 3- to 5-year-old children. *Cognitive Development*, 28, 386-399.

Friston, K. J., Lin, M., Frith, C. D., Pezzulo, G., Hobson, J. A., & Ondobaka, S. (2017). Active Inference, Curiosity and Insight. *Neural Computation*, *29*(10), 2633–2683. https://doi.org/10.1162/neco_a_00999

Frömer, R., & Shenhav, A. (2022). Filling the gaps: Cognitive control as a critical lens for understanding mechanisms of value-based decision-making. *Neuroscience & Biobehavioral Reviews*, *134*, 104483.

Frömer, R., Dean Wolf, C. K., & Shenhav, A. (2019). Goal congruency dominates reward value in accounting for behavioral and neural correlates of value-based decision-making. *Nature communications*, *10*(1), 4926.

Gartus, A., & Leder, H. (2017). Predicting perceived visual complexity of abstract patterns using computational measures: The influence of mirror symmetry on complexity perception. *PloS one*, *12*(11), e0185276.

German, T. P., & Defeyter, M. A. (2000). Immunity to functional fixedness in young children. *Psychonomic Bulletin & Review*, *7*, 707-712.

Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. *Cognition*, *173*, 34-42.

Gluckman, M., & Johnson, S. P. (2013). Attentional capture by social stimuli in young infants. *Frontiers in psychology*, *4*, 527.

Golman, R., & Loewenstein, G. (2018a). The desire for knowledge and wisdom. *The New Science of Curiosity, G. Gordon (ed.), Nova*.

Golman, R., & Loewenstein, G. (2018b). Information gaps: A theory of preferences regarding the presence and absence of information. *Decision*, *5*(3), 143–164. https://doi.org/10.1037/dec0000068

Gopnik, A. (2020). Childhood as a solution to explore–exploit tensions. *Philosophical Transactions of the Royal Society B*, *375*(1803), 20190502.

Gopnik, A., O'Grady, S., Lucas, C. G., Griffiths, T. L., Wente, A., Bridgers, S., ... & Dahl, R. E. (2017). Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proceedings of the National Academy of Sciences*, *114*(30), 7892-7899.

Gori, M., Del Viva, M., Sandini, G., & Burr, D. C. (2008). Young children do not integrate visual and haptic form information. *Current Biology*, *18*(9), 694-698.

Gottlieb, J., & Oudeyer, P.-Y. (2018). Towards a neuroscience of active sampling and curiosity. *Nature Reviews Neuroscience*, *19*(12), 758–770. https://doi.org/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 Sciences*, *17*(11), 585–593. https://doi.org/10.1016/j.tics.2013.09.001

Gruber, M. J., & Ranganath, C. (2019). How curiosity enhances hippocampus-dependent memory: The prediction, appraisal, curiosity, and exploration (PACE) framework. *Trends in cognitive sciences*, *23*(12), 1014-1025.

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

Goupil, L., & Proust, J. (2023). Curiosity as a metacognitive feeling. *Cognition*, *231*, 105325.

Goupil, L., Romand-Monnier, M., & Kouider, S. (2016). Infants ask for help when they know they don't know. *Proceedings of the National Academy of Sciences*, *113*(13), 3492-3496.

Guilbert, J., Jouen, F., & Molina, M. (2018). Motor imagery development and proprioceptive integration: Which sensory reweighting during childhood?. *Journal of experimental child psychology*, *166*, 621-634.

Güçlütürk, Y., Jacobs, R. H., & Lier, R. V. (2016). Liking versus complexity: Decomposing the inverted U-curve. *Frontiers in human neuroscience*, *10*, 112.

Ha, O. R., Bruce, A. S., Pruitt, S. W., Cherry, J. B. C., Smith, T. R., Burkart, D., ... & Lim, S. L. (2016). Healthy eating decisions require efficient dietary self-control in children: A mouse-tracking food decision study. *Appetite*, *105*, 575-581.

Hehman, E., Stolier, R. M., & Freeman, J. B. (2015). Advanced mouse-tracking analytic techniques for enhancing psychological science. *Group Processes & Intergroup Relations*, *18*(3), 384-401.

Henderson, B., & Moore, S. G. (1980). Children's responses to objects differing in novelty in relation to level of curiosity and adult behavior. *Child development*, 457-465.

Hostetter, A. B., & Alibali, M. W. (2019). Gesture as simulated action: Revisiting the framework. *Psychonomic bulletin & review,* 26(3), 721-752.

Hostetter, A. B., & Alibali, M.W. (2008). Visible embodiment: Gestures as simulated action. *Psychonomic Bulletin & Review*, 15, 495–514.

Howard, S. J., & Melhuish, E. (2017). An early years toolbox for assessing early executive function, language, self-regulation, and social development: Validity, reliability, and preliminary norms. *Journal of Psychoeducational Assessment*, *35*(3), 255-275.

Hutt, C., & McGRew, P. L. (1969). Do children really prefer visual complexity?. *Psychonomic Science*, *17*(2), 113-114.

Ikeda, Y., Okuzumi, H., & Kokubun, M. (2014). Age-related trends of inhibitory control in Stroop-like big–small task in 3 to 12-year-old children and young adults. *Frontiers in Psychology*, *5*, 227.

Im-Bolter, N., Johnson, J., & Pascual-Leone, J. (2006). Processing limitations in children with specific language impairment: The role of executive function. *Child development*, *77*(6), 1822-1841.

Jacobsen, T., & Höfel, L. (2001). Aesthetics electrified: An analysis of descriptive symmetry and evaluative aesthetic judgment processes using event-related brain potentials. *Empirical Studies of the Arts*, *19*(2), 177-190.

Jacobsen, T., & Höfel, L. E. A. (2002). Aesthetic judgments of novel graphic patterns: Analyses of individual judgments. *Perceptual and motor skills*, *95*(3), 755-766.

Jakesch, M., & Carbon, C. C. (2012). The mere exposure effect in the domain of haptics. *PloS one*, *7*(2), e31215.

Jepma, M., Schaaf, J. V., Visser, I., & Huizenga, H. M. (2020). Uncertainty-driven regulation of learning and exploration in adolescents: A computational account. *PLoS computational biology*, *16*(9), e1008276.

Jepma, M., Verdonschot, R. G., Van Steenbergen, H., Rombouts, S. A., & Nieuwenhuis, S. (2012). Neural mechanisms underlying the induction and relief of perceptual curiosity. *Frontiers in behavioral neuroscience*, *6*, 5.

Jimenez, S. R. (2018). *Preschoolers' curiosity about novel words*. Vanderbilt University.

Jimenez, S., Sun, Y., & Saylor, M. M. (2018). The process of active word learning. *Active learning from infancy to childhood: Social motivation, cognition, and linguistic mechanisms*, 75-93.

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.

Johnson, S. P. (2004). Development of perceptual completion in infancy. *Psychological Science*, *15*(11), 769-775.

Johnson, S.P., & Aslin, R.N. (1996). Perception of object unity in young infants: The roles of motion, depth, and orientation. *Cognitive Development, 11*, 161-180.

Johnson, M.H., Mareschal, D., & Csibra, G. (2008). The development and integration of the dorsal and ventral visual pathways in object processing. In C.A. Nelson & M. Luciana (Eds.) *The Handbook of Developmental Cognitive Neuroscience (2nd Ed)*, pp. 467-498. Cambridge, MA: MIT press.

Jung, W. P., Kahrs, B. A., & Lockman, J. J. (2015). Manual action, fitting, and spatial planning: Relating objects by young children. *Cognition*, 134, 128-139.

Jung, W. P., Kahrs, B. A., & Lockman, J. J. (2018). Fitting handled objects into apertures by 17- to 36-month-old children: The dynamics of spatial

Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. *Heuristics and biases: The psychology of intuitive judgment*, *49*(49-81), 74.

Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T. Y., & Camerer, C. F. (2009). The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological science*, *20*(8), 963-973.

Karlsson, N., Loewenstein, G., & Seppi, D. (2009). The ostrich effect: Selective attention to information. *Journal of Risk and uncertainty*, *38*, 95-115.

Kashdan, T. B., Stiksma, M. C., Disabato, D. J., McKnight, P. E., Bekier, J., Kaji, J., & Lazarus, R. (2018). The five-dimensional curiosity scale: Capturing the bandwidth of curiosity and identifying four unique subgroups of curious people. *Journal of Research in Personality*, *73*, 130–149. https://doi.org/10.1016/j.jrp.2017.11.011

Kieslich, P. J., & Henninger, F. (2017). Mousetrap: An integrated, open-source mouse-tracking package. *Behavior research methods*, *49*, 1652-1667.

Kieslich, P. J., Schoemann, M., Grage, T., Hepp, J., & Scherbaum, S. (2020). Design factors in mouse-tracking: What makes a difference?. *Behavior Research Methods*, *52*, 317-341.

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

Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks Effect: Human Infants Allocate Attention to Visual Sequences That Are Neither Too Simple Nor Too Complex. *PLoS ONE*, *7*(5), e36399. https://doi.org/10.1371/journal.pone.0036399

Kindlon, D., Mezzacappa, E., & Earls, F. (1995). Psychometric properties of impulsivity measures: Temporal stability, validity and factor structure. *Journal of Child Psychology and Psychiatry*, *36*(4), 645-661.

Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. *Nature human behaviour*, *3*(6), 587-595.

Koop, G. J., & Johnson, J. G. (2011). Response dynamics: A new window on the decision process. *Judgment and Decision making*, *6*(8), 750-758.

Lacey, S., Campbell, C., & Sathian, K. (2007). Vision and touch: multiple or multisensory representations of objects?. *Perception*, *36*(10), 1513-1521.

Lakhal, S., Darmon, A., Bouchaud, J. P., & Benzaquen, M. (2020). Beauty and structural complexity. *Physical Review Research*, *2*(2), 022058.

Lakoff, G., Johnson, M., & Sowa, J. F. (1999). Review of Philosophy in the Flesh: The embodied mind and its challenge to Western thought. *Computational Linguistics*, *25*(4), 631-634.

Lane, S. H. (1968). Preference for complexity as a function of schematic orderliness and redundancy. *Psychonomic Science*, *13*, 209-210

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), e13178.

Lau, J. K. L., Ozono, H., Kuratomi, K., Komiya, A., & Murayama, K. (2020). Shared striatal activity in decisions to satisfy curiosity and hunger at the risk of electric shocks. *Nature Human Behaviour*, *4*(5), 531-543.

Lederman, S. J., & Klatzky, R. L. (1987). Hand movements: A window into haptic object recognition. *Cognitive psychology*, *19*(3), 342-368.

Lee, D. G., & Hare, T. A. (2022). Evidence accumulates for individual attributes during value-based decisions. *Decision*.

Lehman, J., & Stanley, K. O. (2011). Novelty search and the problem with objectives. *Genetic programming theory and practice IX*, 37-56.

Li, R., Roberts, R. C., Huettel, S. A., & Brannon, E. M. (2017). Five-year-olds do not show ambiguity aversion in a risk and ambiguity task with physical objects. *Journal of experimental child psychology*, *159*, 319-326.

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.

Litman, J. (2005). Curiosity and the pleasures of learning: Wanting and liking new information. *Cognition & Emotion*, *19*(6), 793–814. https://doi.org/10.1080/02699930541000101

Litman, J. A. (2008). Interest and deprivation factors of epistemic curiosity. *Personality and individual differences*, *44*(7), 1585-1595.

Litman, J. A., Collins, R. P., & Spielberger, C. D. (2005). The nature and measurement of sensory curiosity. *Personality and Individual Differences*, *39*(6), 1123-1133.

Litman, J. A., & Jimerson, T. L. (2004). The measurement of curiosity as a feeling of deprivation. *Journal of personality assessment*, *82*(2), 147-157.

Liquin, E. G., & Gopnik, A. (2022). Children are more exploratory and learn more than adults in an approach-avoid task. *Cognition*, *218*, 104940.

Liquin, E. G., Callaway, F., & Lombrozo, T. (2021). Developmental change in what elicits curiosity. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 43, No. 43).

Lo, S., & Andrews, S. (2015). To transform or not to transform: Using generalized linear mixed models to analyse reaction time data. *Frontiers in psychology*, *6*, 1171.

Locher, P. J., & Nodine, C. F. (1987). Symmetry catches the eye. In *Eye movements from physiology to cognition* (pp. 353-361). Elsevier.

Locher, P. J., & Simmons, R. W. (1978). Influence of stimulus symmetry and complexity upon haptic scanning strategies during detection, learning, and recognition tasks. *Perception & psychophysics*, *23*, 110-116.

Lockman, J. J. (2000). A perception-action perspective on tool use development. *Child Development*,71, 137-144.

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological bulletin*, *116*(1), 75.

Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin*, *127*(2), 267–286. https://doi.org/10.1037/0033-2909.127.2.267

Lucas, C. G., Bridgers, S., Griffiths, T. L., & Gopnik, A. (2014). When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition*, *131*(2), 284-299.

Luciw, M., Kompella, V., Kazerounian, S., & Schmidhuber, J. (2013). An intrinsic value system for developing multiple invariant representations with incremental slowness learning. *Frontiers in neurorobotics*, *7*, 9.

Madan, C. R., Spetch, M. L., & Ludvig, E. A. (2015). Rapid makes risky: Time pressure increases risk seeking in decisions from experience. *Journal of Cognitive Psychology*, *27*(8), 921-928.

Maldonado, M., Dunbar, E., & Chemla, E. (2019). Mouse tracking as a window into decision making. *Behavior research methods*, *51*, 1085-1101.

Marghetis, T., Núñez, R., & Bergen, B. K. (2014). Doing arithmetic by hand: Hand movements during exact arithmetic reveal systematic, dynamic spatial processing. *Quarterly Journal of Experimental Psychology*, *67*(8), 1579-1596.

Marin, M. M., & Leder, H. (2018). Exploring aesthetic experiences of females: Affect-related traits predict complexity and arousal responses to music and affective pictures. *Personality and individual Differences*, *125*, 80-90

Martindale, C., Moore, K., & Borkum, J. (1990). Aesthetic preference: Anomalous findings for Berlyne's psychobiological theory. *The American Journal of Psychology*, 53-80.

Matta, A. D., Gonçalves, F. L., & Bizarro, L. (2012). Delay discounting: Concepts and measures. *Psychology & Neuroscience*, *5*, 135-146.

Mcdougall, S. J., Curry, M. B., & De Bruijn, O. (1999). Measuring symbol and icon characteristics: Norms for concreteness, complexity, meaningfulness, familiarity, and semantic distance for 239 symbols. *Behavior Research Methods, Instruments, & Computers*, *31*(3), 487-519.

Meder, B., Wu, C. M., Schulz, E., & Ruggeri, A. (2021). Development of directed and random exploration in children. *Developmental science*, *24*(4), e13095.

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., ... & Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, *2*(3), 191.

Messinger, S. M. (1998). Pleasure and complexity: Berlyne revisited. *The Journal of Psychology*, *132*(5), 558-560.

Metcalfe, J., Schwartz, B. L., & Eich, T. S. (2020). Epistemic curiosity and the region of proximal learning. *Current Opinion in Behavioral Sciences*, *35*, 40–47. https://doi.org/10.1016/j.cobeha.2020.06.007

Minuchin, P. (1971). Correlates of curiosity and exploratory behavior in preschool disadvantaged children. *Child Development*, 939-950.

Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgment and Decision making*, *5*(6), 437-449.

Molinaro, G., & Collins, A. G. (2023). A goal-centric outlook on learning. *Trends in Cognitive Sciences*.

Munsinger, H., & Kessen, W. (1966). Structure, variability, and development. *Journal of Experimental Child Psychology*, *4*(1), 20-49.

Munsinger, H., Kessen, W., & Kessen, M. L. (1964). Age and uncertainty: Developmental variation in preference for variability. *Journal of Experimental Child Psychology*, *1*(1), 1-15.

Murayama, K. (2022). A reward-learning framework of knowledge acquisition: An integrated account of curiosity, interest, and intrinsic–extrinsic rewards. *Psychological Review*, *129*(1), 175.

Murayama, K., FitzGibbon, L., & Sakaki, M. (2019). Process Account of Curiosity and Interest: A Reward-Learning Perspective. *Educational Psychology Review*, *31*(4), 875–895. https://doi.org/10.1007/s10648-019-09499-9

Muth, C., Ebert, S., Marković, S., & Carbon, C. C. (2019). "Aha" ptics: enjoying an aesthetic aha during haptic exploration. *Perception*, *48*(1), 3-25.

Nadal, M., Munar, E., Marty, G., & Cela-Conde, C. J. (2010). Visual complexity and beauty appreciation: Explaining the divergence of results. *Empirical Studies of the*

Arts, *28*(2), 173-191.

Nakamura, J., & Csikszentmihalyi, M. (2002). The concept of flow. *Handbook of positive psychology*, *89*, 105.

Nath, S. S., Brändle, F., Schulz, E., Dayan, P., & Brielmann, A. A. (2023). Relating Objective Complexity, Subjective Complexity and Beauty.

Nussenbaum, K., Velez, J. A., Washington, B. T., Hamling, H. E., & Hartley, C. A. (2022). Flexibility in valenced reinforcement learning computations across development. *Child development*, *93*(5), 1601-1615.

O'Hora, D., Carey, R., Kervick, A., Crowley, D., & Dabrowski, M. (2016). Decisions in motion: Decision dynamics during intertemporal choice reflect subjective evaluation of delayed rewards. *Scientific reports*, *6*(1), 20740.

Olschewski, S., & Rieskamp, J. (2021). Distinguishing three effects of time pressure on risk taking: Choice consistency, risk preference, and strategy selection. *Journal of Behavioral Decision Making*, *34*(4), 541-554.

O'Neill, G., & Miller, P. H. (2013). A show of hands: Relations between young children's gesturing and executive function. *Developmental psychology*, *49*(8), 1517. 15

Ornkloo, H., & von Hofsten, C. (2007). Fitting objects into holes: On the development of spatial cognition skills. *Developmental Psychology*, 43, 404-416.

Ornkloo, H., & von Hofsten, C. (2009). Young children's ability to solve spatial problem involving a choice. *European Journal of Developmental Psychology*, 6, 685-704.

Ossmy, O., Han, D., Cheng, M., Kaplan, B. E., & Adolph, K. E. (2020). Look before you fit: The real-time planning cascade in children and adults. *Journal of Experimental Child Psychology*, 189, 104-696.

Otto, A. R., Knox, W. B., Markman, A. B., & Love, B. C. (2014). Physiological and behavioral signatures of reflective exploratory choice. *Cognitive, Affective, & Behavioral Neuroscience*, *14*, 1167-1183.

Oudeyer, P.-Y., Gottlieb, J., & Lopes, M. (2016). Intrinsic motivation, curiosity, and learning. In *Progress in Brain Research* (Vol. 229, pp. 257–284). Elsevier. https://doi.org/10.1016/bs.pbr.2016.05.005

Overvliet, K. E., Krampe, R. T., & Wagemans, J. (2012). Perceptual grouping in haptic search: The influence of proximity, similarity, and good continuation. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(4), 817.

Overvliet, K. E., & Sayim, B. (2016). Perceptual grouping determines haptic contextual modulation. *Vision Research*, *126*, 52-58.

Owen, D. H., & Brown, D. R. (1970). Visual and tactual form discrimination: Psychophysical comparison within and between modalities. *Perception & Psychophysics*, *7*(5), 302-306.

Pearce, A. L., Adise, S., Roberts, N. J., White, C., Geier, C. F., & Keller, K. L. (2020). Individual differences in the influence of taste and health impact successful dietary self-control: A mouse tracking food choice study in children. *Physiology & behavior*, *223*, 112990. https://doi.org/10.1016/j.physbeh.2020.112990

Plebanek, D. J., & Sloutsky, V. M. (2017). Costs of selective attention: When children notice what adults miss. *Psychological science*, *28*(6), 723-732.

Poli, F., Meyer, M., Mars, R. B., & Hunnius, S. (2022). Contributions of expected learning progress and perceptual novelty to curiosity-driven exploration. *Cognition*, *225*, 105119. https://doi.org/10.1016/j.cognition.2022.105119

Poli, F., Serino, G., Mars, R. B., & Hunnius, S. (2020). Infants tailor their attention to maximize learning. *Science Advances*, *6*(39), eabb5053. https://doi.org/10.1126/sciadv.abb5053

Pouw, W. T., De Nooijer, J. A., Van Gog, T., Zwaan, R. A., & Paas, F. (2014). Toward a more embedded/extended perspective on the cognitive function of gestures. *Frontiers in psychology*, *5*, 359.

R Core Team (2022). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL https://www.R-project.org/

Rigney, J., & Wang, S. H. (2015). Delineating the boundaries of infants' spatial categories: The case of containment. *Journal of Cognition and Development*, 16, 420-441

Rogers, R. D., & Monsell, S. (1995). Costs of a predictible switch between simple cognitive tasks. *Journal of experimental psychology: General*, *124*(2), 207.

Rosenbaum, G. M., & Hartley, C. A. (2019). Developmental perspectives on risky and impulsive choice. *Philosophical Transactions of the Royal Society B*, *374*(1766), 20180133.

Ryan, R. M., & Deci, E. L. (2000). Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary educational psychology*, *25*(1), 54-67.

Ryan, R. M., & Deci, E. L. (2017). Self-determination theory. *Basic psychological needs in motivation, development, and wellness*.

Ruggeri, A., Pelz, M., Gopnik, A, & Schulz, E. (2021). *Toddlers search longer when there is more information to be gained* [Preprint]. PsyArXiv. https://doi.org/10.31234/osf.io/uzdvp

Rump, E. E. (1968). Relative preference as a function of the number of elements in abstract designs. *Australian Journal of Psychology*, *20*(1), 39-48.

Sadeghiyeh, H., Wang, S., Alberhasky, M. R., Kyllo, H. M., Shenhav, A., & Wilson, R. C. (2020). Temporal discounting correlates with directed exploration but not with random exploration. *Scientific reports*, *10*(1), 4020.

Scherbaum, S., & Kieslich, P. J. (2018). Stuck at the starting line: How the starting procedure influences mouse-tracking data. *Behavior Research Methods*, *50*, 2097-2110.

Schoemann, M., Lüken, M., Grage, T., Kieslich, P. J., & Scherbaum, S. (2019). Validating mouse-tracking: How design factors influence action dynamics in intertemporal decision making. *Behavior Research Methods*, *51*, 2356-2377.

Schröer, L., Cooper, R. P., & Mareschal, D. (2021). Science with Duplo: Multilevel goal management in preschoolers' toy house constructions. *Journal of Experimental Child Psychology*, *206*, 105067.

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

Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for rewards like a child means less generalization and more directed exploration. *Psychological science*, *30*(11), 1561-1572.

Schwartenbeck, P., Passecker, J., Hauser, T. U., FitzGerald, T. H., Kronbichler, M., & Friston, K. J. (2019). Computational mechanisms of curiosity and goal-directed exploration. *eLife*, *8*, e41703. https://doi.org/10.7554/eLife.41703

Silvia, P. J. (2005). What is interesting? Exploring the appraisal structure of interest. *Emotion*, *5*(1), 89.

Sim, Z. L., & Xu, F. (2017). Infants preferentially approach and explore the unexpected. *British Journal of Developmental Psychology*, *35*(4), 596-608.

Smith, L., & Gasser, M. (2005). The development of embodied cognition: Six lessons from babies. *Artificial life*, *11*(1-2), 13-29.

Smock, C. D., & Holt, B. G. (1962). Children's Reactions to Novelty: An Experimental Study of" Curiosity Motivation". *Child Development*, 631-642.

Sobeh, J., & Spijkers, W. (2013). Development of neuropsychological functions of attention in two cultures: A cross-cultural study of attentional performances of Syrian and German children of pre-school and school age. *European Journal of Developmental Psychology*, *10*(3), 318-336.

Somerville, L. H., & Casey, B. J. (2010). Developmental neurobiology of cognitive control and motivational systems. *Current opinion in neurobiology*, *20*(2), 236-241.

Somerville, L. H., Sasse, S. F., Garrad, M. C., Drysdale, A. T., Abi Akar, N., Insel, C., & Wilson, R. C. (2017). Charting the expansion of strategic exploratory behavior during adolescence. *Journal of experimental psychology: general*, *146*(2), 155.

Song, J. H., & Nakayama, K. (2006). Role of focal attention on latencies and trajectories of visually guided manual pointing. *Journal of vision*, *6*(9), 11-11.

Soska, K. C., & Johnson, S. P. (2008). Development of three-dimensional object completion in infancy. *Child Development*, 79, 1230-1236.

Spencer, J. P., Samuelson, L. K., Blumberg, M. S., McMurray, B., Robinson, S. R., & Bruce Tomblin, J. (2009). Seeing the world through a third eye: Developmental systems theory looks beyond the nativist–empiricist debate. *Child development perspectives*, *3*(2), 103-105.

Spielberger, C. D., Peters, R. A., & Frain, F. (1976). The state-trait curiosity inventory. *Unpublished manual, University of South Florida, Tampa*.

Spielberger, C. D., & Starr, L. M. (1994). Curiosity and exploratory behavior. In H. F. O'Neil, Jr. & M. Drillings (Eds.), *Motivation: Theory and research* (pp. 221±243). Hillsdale, NJ: Erlbaum.

Spivey, M. J., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences*, *102*(29), 10393-10398.

Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, *348*(6230), 91-94.

Street, S. Y., James, K. H., Jones, S. S., & Smith, L. B. (2011). Vision for action in toddlers: The posting task. *Child development*, *82*(6), 2083-2094.

Sullivan, N., Hutcherson, C., Harris, A., & Rangel, A. (2015). Dietary self-control is related to the speed with which attributes of healthfulness and tastiness are processed. *Psychological science*, *26*(2), 122-134.

Sumner, E., Li, A. X., Perfors, A., Hayes, B., Navarro, D., & Sarnecka, B. W. (2019). The Exploration Advantage: Children's instinct to explore allows them to find information that adults miss. *PsyArxiv*. *https://psyarxiv.com/h437v/download?format=pdf*

Sumner, E. S., Steyvers, M., & Sarnecka, B. W. (2019). It's not the treasure, it's the hunt: Children are more explorative on an explore/exploit task than adults. In *CogSci* (pp. 2891-2897).

Sun, Z., & Firestone, C. (2021). Curious objects: How visual complexity guides attention and engagement. *Cognitive Science,* 45(4), e12933.

Sun, Z., & Firestone, C. (2022). Beautiful on the inside: Aesthetic preferences and the skeletal complexity of shapes. *Perception,* 51(12), 904-918.

Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning* (Vol. 135, pp. 223-260). Cambridge: MIT press.

Taffoni, F., Tamilia, E., Focaroli, V., Formica, D., Ricci, L., Di Pino, G., ... & Keller, F. (2014). Development of goal-directed action selection guided by intrinsic motivations: an experiment with children. *Experimental brain research*, *232*, 2167-2177.

Tajima, S., Drugowitsch, J., & Pouget, A. (2016). Optimal policy for value-based decision-making. *Nature communications*, *7*(1), 12400.

Tang, H., Houthooft, R., Foote, D., Stooke, A., Chen, O. X., Duan, Y., ... Abbeel, P. (2017). Exploration: A study of count-based exploration for deep reinforcement learning. In *Advances in Neural Information Processing Systems* (pp. 2750–2759)

Ten, A., Kaushik, P., Oudeyer, P.-Y., & Gottlieb, J. (2021). Humans monitor learning progress in curiosity-driven exploration. *Nature Communications*, *12*(1), Article 1. https://doi.org/10.1038/s41467-021-26196-w

Thompson-Schill, S. L., Ramscar, M., & Chrysikou, E. G. (2009). Cognition Without Control: When a Little Frontal Lobe Goes a Long Way. *Current Directions in Psychological Science*, *18*(5), 259–263. https://doi.org/10.1111/j.1467-8721.2009.01648.x

Tomlinson Jr, J. M., Bailey, T. M., & Bott, L. (2013). Possibly all of that and then some: Scalar implicatures are understood in two steps. *Journal of memory and language*, *69*(1), 18-35.

Tomov, M. S., Truong, V. Q., Hundia, R. A., & Gershman, S. J. (2020). Dissociable neural correlates of uncertainty underlie different exploration strategies. *Nature Communications*, *11*(1), 2371. https://doi.org/10.1038/s41467-020-15766-z

Travers, E., Rolison, J. J., & Feeney, A. (2016). The time course of conflict on the Cognitive Reflection Test. *Cognition*, *150*, 109-118.

Twomey, K. E., & Westermann, G. (2018). Curiosity-based learning in infants: A neurocomputational approach. *Developmental Science*, *21*(4), e12629. https://doi.org/10.1111/desc.12629 van Lieshout, L. L., de Lange, F. P., & Cools, R. (2021). Uncertainty increases curiosity, but decreases happiness. *Scientific Reports*, *11*(1), 14014.

van Lieshout, L. L., Traast, I. J., de Lange, F. P., & Cools, R. (2021). Curiosity or savouring? Information seeking is modulated by both uncertainty and valence. *PloS one*, *16*(9), e0257011.

van Schijndel, T. J., Visser, I., van Bers, B. M., & Raijmakers, M. E. (2015). Preschoolers perform more informative experiments after observing theory-violating evidence. *Journal of experimental child psychology*, *131*, 104-119.

Vellani, V., de Vries, L. P., Gaule, A., & Sharot, T. (2020). A selective effect of dopamine on information-seeking. *Elife*, *9*, e59152.

Vogl, E., Pekrun, R., Murayama, K., & Loderer, K. (2020). Surprised–curious–confused: Epistemic emotions and knowledge exploration. *Emotion*, *20*(4), 625–641. https://doi.org/10.1037/emo0000578

Vygotsky, L. S. (1978). The role of play in development.

Wakefield, E. M., Novack, M. A., & Goldin-Meadow, S. (2018). Unpacking the ontogeny of gesture understanding: How movement becomes meaningful across development. *Child development*, *89*(3), e245-e260.

Welsh, T. N. (2011). The relationship between attentional capture and deviations in movement trajectories in a selective reaching task. *Acta psychologica*, *137*(3), 300-308.

Wilson, R. C., Bonawitz, E., Costa, V. D., & Ebitz, R. B. (2021). Balancing exploration and exploitation with information and randomization. *Current Opinion in Behavioral Sciences*, *38*, 49–56. https://doi.org/10.1016/j.cobeha.2020.10.001

Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, *143*(6), 2074.

Wirth, R., Foerster, A., Kunde, W., & Pfister, R. (2020). Design choices: Empirical recommendations for designing two-dimensional finger-tracking experiments. *Behavior Research Methods*, *52*, 2394-2416.

Witten, I. H. (1977). Exploring, modelling and controlling discrete sequential environments. *International Journal of Man-Machine Studies*, *9*(6), 715-735.

Wohlwill, J. F. (1968). Amount of stimulus exploration and preference as differential functions of stimulus complexity. *Perception & Psychophysics*, *4*(5), 307-312.

Woolley, K., & Fishbach, A. (2017). Immediate rewards predict adherence to long-term goals. *Personality and Social Psychology Bulletin*, *43*(2), 151-162.

Wu, S., Blanchard, T., Meschke, E., Aslin, R. N., Hayden, B. Y., & Kidd, C. (2022). Macaques preferentially attend to intermediately surprising information. *Biology letters*, *18*(7), 20220144.

Wu, R., Gopnik, A., Richardson, D. C., & Kirkham, N. Z. (2011). Infants learn about objects from statistics and people. *Developmental psychology*, *47*(5), 1220.

Wu, C. M., Schulz, E., Pleskac, T. J., & Speekenbrink, M. (2022). Time pressure changes how people explore and respond to uncertainty. *Scientific reports*, *12*(1), 4122.

Wulff, D. U., Haslbeck, J. M., Kieslich, P. J., Henninger, F., & Schulte-Mecklenbeck, M. (2019). Mouse-tracking: Detecting types in movement trajectories. In *A handbook of process tracing methods* (pp. 131-145). Routledge.

Wulff, D. U., Kieslich, P. J., Henninger, F., Haslbeck, J. M., & Schulte-Mecklenbeck, M. (2021). Movement tracking of cognitive processes: A tutorial using mousetrap. *PsyArXiv*, *10*.

Zelazo, P. D. (2006). The Dimensional Change Card Sort (DCCS): A method of assessing executive function in children. *Nature protocols*, *1*(1), 297-301.

Zelazo, P. D., Anderson, J. E., Richler, J., Wallner-Allen, K., Beaumont, J. L., & Weintraub, S. (2013). II. NIH Toolbox Cognition Battery (CB): Measuring executive function and attention. *Monographs of the Society for Research in Child Development*, *78*(4), 16-33.