UNIVERSITY of STIRLING

Developing Distinctively Human Cumulative Culture:

Age-Related Changes in Social Information Use

Kirsten Helen Blakey

University of Stirling

Submitted for the degree of Doctor of Philosophy.

December 2020

Thesis Abstract

This thesis investigated the distinctiveness of human cumulative culture by examining the developmental trajectory of reasoning-based social learning strategies, which have been proposed to be what sets human learning apart from non-humans. Specifically, the studies reported in this thesis were concerned with differentiating cases in which social information use was driven by reasoned understanding and cases which could be explained by implicit adaptive heuristics. This was achieved by looking for age-related changes in children's reasoning about, and use of, social information. More effective social information use was proposed to reflect learners' reasoned understanding of its relevance and potential value to themselves. Each study examined a particular cognitive challenge identified as potentially relevant for social information use in the context of real world cases of cumulative culture. Chapter two explored the development of children's ability to account for others' conflicting goals in their use of the available social information as a means to achieve their own goal. Chapters three and four investigated children's ability to seek out appropriate sources of social information. Chapter three looked at children's recognition of what information they required to solve a problem and who could provide that information. While chapter four examined children's ability to consider potential informants' mental states when determining 'who knows'. Overall, the developmental trajectory indicated relatively late childhood development of effective social information use driven by reasoned understanding. This late development is consistent with proposals suggesting that this may be a cognitive mechanism that is only available to humans. The flexibility afforded by the ability to recognise the value, to oneself, of others' potential to provide useful and relevant information, on account of their experience or knowledge, appears to offer the significant advantage in social information use that may drive human cumulative culture beyond the capabilities of non-humans.

Acknowledgements

My first thanks must go to all of the children who participated in my research (and of course their parents/guardians). I immensely enjoyed working with them all and seeing them enjoy taking part made the months of planning, designing, and creating task materials more than worth it. A special thanks goes to RZSS Edinburgh Zoo, specifically the staff at the Living Links and Budongo Research Unit who warmly welcomed me to collect data for all of my studies. Being able to watch the chimpanzees in between testing participants provided great entertainment. I am also grateful for the support of the staff in all of the schools and nurseries who made me feel so welcome. Particular thanks must also go to the staff and children in the Psychology Kindergarten at the University of Stirling who helped so much with piloting studies and acting as models for the study reported in Chapter three.

I thank the Division of Psychology at the University of Stirling and the European Research Council for funding my PhD studentship which made it possible for me to conduct my research as well as present it at some amazing conferences.

Most importantly I would like to thank my supervisor, Prof Christine Caldwell, whose guidance, encouragement, feedback, and continual reassurance has made this PhD a really fantastic learning experience. I admire her knowledge, experience, and commitment to great research. I also want to thank my second supervisor, Dr Eva Rafetseder, for her unwavering support and enthusiasm. Though they were not my supervisors, I am also incredibly grateful to Dr Liz Renner and Dr Mark Atkinson for all of their advice. Specifically, their encouragement to learn to code in both R and PsychoPy has given me the confidence to push myself to learn new skills.

I am really grateful to have shared my PhD experience with best and most supportive PhD buddies, Gemma Mackintosh, Juliet Robertstone, Donna Kean, and Charlotte Wilks. I thank them for excellent discussions, dealing with my incessant questions, helping me make task materials, allowing me to talk through data analysis, incredible emotional support, and for their friendship.

Completing a PhD thesis during the COVID-19 pandemic has certainly been a challenge. However, it would have been much harder without the hugely supportive online community of academic writers that I found in the VWR-PGR virtual writing retreat. I’m sure that I could not have maintained my motivation to write without them.

Finally, I want to thank my family. I thank my parents, Colin and Sheena, and my brother, Thomas, for accepting that eight years at university is not too long, for humouring my constant phone calls to give updates on my research, and for loving and supporting me no matter what. I would also like to mention my grandparents: Gran, who would love to have been able to do this research in my place; Grandpa, who doesn’t really understand what I do, but loves to ask questions anyway; and Grandma, who saw the start but not the end of this adventure.

Contents

Chapter 1:

General Introduction

Unquestionably there are striking differences between humans and non-humans. To understand our peculiarity and what drives the success of our species, is really to ask what it means to be human. Humans' monopoly of the planet is proposed to be enabled by our apparently unique ability to exploit and build on the accomplishments of previous generations (cumulative culture; Boyd & Richerson, 1996; Henrich & McElreath, 2003). That is, the resultant technologies, skills, and knowledge of later generations offer greater benefits relative to what could be achieved by an individual (Tomasello et al., 1993). By contrast, claims of comparable abilities in non-humans are rare and usually contentious. Unsurprisingly, questions regarding the origins, development, and drivers of human uniqueness fascinate researchers across a wealth of disciplines. Indeed, they have been posed as some of the most important and yet unanswered scientific questions ("So much more to know," 2005).

The stark contrast in the scope of cumulative culture in human adults compared to the limited evidence in non-humans has led to proposals of a uniquely human cognitive mechanism that is a necessary prerequisite for human-like cumulative culture. One way to approach investigation of such cognitive mechanisms is to examine the developmental trajectory of cumulative culture in human children. In this thesis I investigate the proposal that human cumulative culture relies on the optimisation of social information use driven by capacities that enable understanding of the value and relevance of social information. Looking at whether age-related changes in social information use coincide with advances in cognitive development could offer an insight into the underlying mechanisms that are required, and help to explain why cumulative culture appears to be restricted to humans. Therefore, this thesis aims to establish the cognitive mechanisms that drive the development of distinctively human cumulative culture by tackling fundamental questions related to the developmental trajectory of cognitive capacities for social information use.

In this chapter I will review the background literature that informed the rationale for the experimental work undertaken in chapters two, three, and four, outlining the specific aims and hypotheses being examined. I begin by outlining the current state of the literature on cumulative culture including: a definition, a review of evidence in human adults relative to non-humans, and an outline of the potential value of developmental research in understanding the discontinuity between these populations and establishing prerequisite sociocognitive mechanisms. I will then move onto reviewing the sociocognitive mechanisms that have been proposed to underlie this seemingly human unique propensity, highlighting capacities that appear to be shared by humans and non-humans, and those that appear to differ between species. Specifically, I will outline how looking at children's use of social information, and their awareness of doing so, could offer an insight into their understanding of what makes that social information valuable. Understanding the value of social information is vital, as it is likely important for successful cumulative culture across a broad range of contexts. There is little evidence of distinctive adult-like cumulative culture in children and non-humans. Therefore, examining social information use in these populations could provide clues as to their capacity for cumulative culture and the sociocognitive mechanisms that might be necessary for it. I will focus in particular on a recent proposal that suggests that distinctively human cumulative culture is driven by a particular type of strategic social learning that is available only to humans.

In another section of this chapter I will deviate from the discussion of cumulative culture and focus instead on children's cognitive development. As cognitive development significantly informs the aims and hypotheses of this thesis it is pertinent to provide sufficient background. Therefore, I will review the ontogeny of cognitive development, providing evidence for the continuing advancement of cognitive capacities throughout development from infancy to adolescence. In doing so I will highlight key milestones in cognitive development. Following this I will set out the proposed relationship between children's cognitive development and the development of distinctively human cumulative culture, laying the foundation for the aims and hypotheses of this thesis.

Distinctively Human Cumulative Culture

Social transmission can result in an accumulation of beneficial modifications to cultural traits over successive generations of learners. Usually referred to as cumulative cultural evolution, or cumulative culture (Caldwell & Millen, 2008b), this directional pattern of change has also been aptly dubbed the 'ratchet effect' (Tomasello et al., 1993) on account of the preservation of the modifications to cultural traits. The consequent advances in knowledge and skills are such that the cultural traits of later generations, afforded by their access to social information, are relatively unambiguous 'improvements' in that they would likely have been preferred by their predecessors had they been available (Caldwell et al., 2020). Unsurprisingly there have been a variety of definitions of such terms and discontinuity in their uses between disciplines and researchers. Recently Mesoudi and Thornton (2018) identified four core criteria that are widely agreed upon as indicative of cumulative culture: (1) a change in behaviour or cultural product, (2) social transmission of the modified trait, (3) improvement in performance as a result of the modified trait, and (4) iteration of these steps generating improvement over time. Although some researchers include additional constraints for traits to be considered as cumulative culture (e.g., Derex et al., 2018; Enquist et al., 2011; Lewis & Laland, 2012; Mesoudi & Thornton refer to these as 'extended criteria') the core criteria are common across definitions.

These criteria clearly distinguish cumulative culture from cultural evolution in the broader sense. Cultural evolution refers to the process of change over time in the behaviours exhibited by a population as a result of social transmission. Whilst novel modifications of existing cultural traits will show differential success in their rates of transmission and adoption, group-level changes are not necessarily characterised by a cumulative improvement over generations of learners. That is, cultural evolution is akin to iteration of the first two steps of Mesoudi & Thornton's core criteria.

Human populations demonstrate cumulative culture across a breadth of domains (Kempe et al., 2014), including technological processes (e.g., the manufacture of stone tools, Stout, 2011, and communication systems, Smaldino & Richerson, 2013) and abstract conceptual knowledge (e.g., science and mathematics, Arbesman, 2011; May, 1966). Indeed, its ubiquity, in conjunction with a lack of evidence in non-humans, has prompted some to claim that cumulative culture is unique to humans (e.g., Dean et al., 2012, 2014; Tennie et al., 2009). Moreover, some researchers have cited *cumulative* culture as being what makes human culture particularly 'special' (Boyd & Richerson, 1996; Heyes, 2012; Tomasello et al., 1993). Certainly, there have been markedly fewer claims of cumulative culture in non-humans (noteworthy exceptions include New Caledonian crows, Hunt & Gray, 2003; and chimpanzees, Sanz et al., 2009) despite widespread evidence of culture and cultural evolution (Allen, 2019; Aplin, 2019; Whiten, 2017a). For example, Hobaiter et al. (2014) found that wild chimpanzees' tool use can be interpreted as cultural given that their behavioural diversity was driven by social transmission. Other species found to demonstrate culture include: birds (e.g., great tits, Aplin et al., 2015), cetaceans (e.g., humpback whales, Allen et al., 2013), and insects (e.g., bumblebees, Alem et al., 2016; fruit flies, Danchin et al., 2018). In regard to cultural evolution, the most compelling examples come from bird and whale song (e.g., yellowhammers, Pipek et al., 2018; humpback whales, Garland et al., 2011; for a comparative review see Garland & McGregor, 2020).

While it is relatively straightforward to identify culture and cultural evolution in natural populations, it is significantly more challenging to find evidence for cumulative culture. By definition, cumulative culture develops slowly over a long time period (multiple generations of transmission, which in naturally occurring traditions may equate to multiple biological generations). Thus, historical data are required in order to objectively evaluate cultural artefacts and behaviours in terms of their relative value or benefits over time. This presents somewhat of a problem for research into non-human cumulative culture. While some human artefacts have survived the test of time, the same cannot be said for artefacts used by non-humans. Recently however, two cases of naturalistic cumulative culture in non-humans have been described, using data drawn from large historical datasets. Examining the migratory behaviour of bighorn sheep and moose, Jesmer et al. (2018) concluded that translocated populations took multiple generations to develop migratory patterns as effective as those of well-established historical populations. This suggested that later generations had benefitted from the accumulated experience of previous generations. In another example, Schofield et al. (2018) proposed that researchers' field observations of Japanese macaques reflected improvements in novel food processing techniques across the 60 years of available data. Analyses of these historical data sets (Jesmer et al., 2018; Schofield et al., 2018) appear to show patterns of gradual improvement that are suggestive of naturally occurring cumulative culture. An experimental study with non-humans has also fulfilled the criteria for cumulative culture, although this too focused on naturally occurring tendencies. Sasaki and Biro (2017) compared homing pigeons' route efficiency by taking advantage of their natural tendency to fly alongside a partner (or flock) when released together. However, it is important to note that these findings with non-humans are very likely to be context-specific (e.g., evidence of cumulative culture is likely to be restricted to route efficiency in homing pigeons or migratory patterns in ungulates).

Research with human adults has successfully simulated cumulative culture in a laboratory setting using generational replacement (or 'microsociety') designs (Caldwell & Millen, 2008a). By adapting these methodologies, further studies aimed to establish the factors that facilitate or inhibit cumulative culture (Caldwell et al., 2016; Caldwell & Millen, 2009; Derex et al., 2019; Fay et al., 2019; Muthukrishna et al., 2014; Zwirner & Thornton, 2015), as a way to identify the prerequisite cognitive capacities. These kinds of studies do show that particular sociocognitive capacities, for example teaching (Caldwell et al., 2018; Zwirner & Thornton, 2015), may aid the transmission of cultural traits. However, they do not wholly explain the distinctiveness of human cumulative culture or the lack, or extreme rarity, of cumulative culture in non-humans seeing as some such capacities are not exclusive to humans. I will return to the issue of common sociocognitive capacities and what this might mean for explaining distinctively human cumulative culture later in this chapter.

The seemingly unbounded nature of cumulative culture in human adults stands in stark contrast to the extremely restricted evidence in non-humans. Even though recent evidence has to some extent quashed prior claims (Dean et al., 2014) of cumulative culture being precluded in non-humans, there is a distinct difference in its expression. Thus, while non-humans may have some capacity for cumulative culture it appears to be restricted relative to the cumulative culture demonstrated in human adults. The challenge then is to determine whether there is a particular reason, from a cognitive perspective, for the discontinuity between humans and non-humans. One approach to bridging the gap between these two populations is to investigate the development of the capacity for cumulative culture in young children.

There are a number of reasons for examining the developmental trajectory of cumulative culture. First, it can be assumed that, as adults are capable of cumulative culture, it likely develops during childhood. Therefore, testing young children's capabilities for cumulative culture across different ages could provide a valuable insight into its ontogeny. It might also provide evidence as to the degree of such abilities. Specifically, whether children demonstrate patterns of cumulative culture that are comparable to human adults, or whether such capacities appear to be restricted similarly to patterns observed in non-humans. Tracking its developmental trajectory may highlight particular ages at which children make significant advances in their capacity for cumulative culture. If these advances occur at a similar age to the development of particular cognitive capacities, then such capacities may be necessary prerequisites. That is, investigating age-related changes in the capacity for cumulative culture could expose, more clearly than studies with human adults, cognitive mechanisms that underlie distinctively human cumulative culture. Such capacities may be beyond the capabilities of non-humans and therefore may help to explain the discontinuity with regard to human-like cumulative culture. Finally, to some degree research in adults is limited given that there is only so much that can be manipulated experimentally. That is, while it is possible to manipulate the availability, type, and source of information, the availability of cognitive mechanisms that can be employed to use the information cannot be manipulated. Therefore, if the aim is to establish cognitive prerequisites for cumulative culture, manipulating learning conditions may not be sufficient to uncover the mechanisms responsible. If for instance the key distinction between human and non-human social learning lies in a capacity for deeper cognitive understanding, rather than a particular type of information processing, it is unlikely that it would be possible to remove or sufficiently restrict this kind of capacity. However, studies of children might enable us to get at this, albeit only through correlational relationships, as younger children may not yet have developed such mechanisms.

While the majority of research in adults has focused on identifying factors that facilitate or inhibit cumulative culture, the question remains as to whether capacities for cumulative culture are even available in young children. Crucially, the age-related development of these capacities, which likely offers the most value in terms of establishing prerequisite cognitive mechanisms, has not really been addressed at all.

A relatively small number of studies claim to have found evidence of cumulative culture in children. For example, Dean et al. (2012) found that, in contrast to chimpanzees and capuchin monkeys, 3- to 4-year-old children's success in reaching the most complex solution in a multistage puzzlebox task was dependent on social support from other group members. Specifically, transmission of solutions correlated with both imitation and teaching measures. Having used a similar method, McGuigan et al. (2017) described that simpler solutions preceded more advanced and rewarding solutions, with particular techniques more prevalent within rather than between groups, thus suggesting a role of social transmission. However, as these studies employ a closed group design this limits the conclusions that can be drawn. For example, in regard to the Dean et al. (2012) study, it is not possible to determine whether or not children's improved performance was explained by their increasing exposure to the task. Indeed, a recent replication, that included a previously absent asocial control condition (Reindl et al., 2020) found that children were capable of reaching the most complex solution without social learning. Thus, closed group studies are not sufficient for concluding the age from which capacities for cumulative culture are available. Consequently, they do not move us further forward in making use of children's development as a means to establish prerequisite cognitive mechanisms. Although these studies offer an insight into children's capacity to use social information, through adopting others' solutions, they provide only limited insights into the potential for accumulating improvements over time.

Generational replacement studies with young children are rare, and those that have been conducted do not provide evidence of the characteristic accumulation of improvements over generations that is seen in adult populations (e.g., Caldwell & Millen, 2008a). For example, Reindl and Tennie (2018) found that 4- to 5-year-old children in later generations did not perform better than earlier generations. Similarly, Tennie et al. (2014) found that exposure to social information did not enable 4-year-olds to produce more efficient solutions. Both Flynn (2008) and Tennie et al. (2014) found that if chains were deliberately initiated with a demonstration of an artificially inefficient method, some children removed redundant elements to implement and transmit more effective solutions resulting in later generations being more successful than earlier generations. However, cumulative culture should result in higher performance following exposure to social information than could be achieved without any social information, rather than just a drift back to a baseline asocial learning level of success. Therefore, the results of these studies cannot really be considered as evidence for cumulative culture.

Finally, some very recent research (Wilks et al., 2021) using a novel methodology, which simulates a transmission chain from individual participant data (Caldwell et al., 2020), has tackled as yet outstanding questions regarding the developmental trajectory of cumulative culture in children. Indeed, this may be the first study that truly addresses the developmental trajetory of cumulative culture. The results of this study indicate that the presence or otherwise of cumulative culture can sometimes depend on information processing demands. That is, when information was presented such that it did not require storage (and indeed would have been difficult to ignore), even very young children demonstrated a ratchet-like effect in their learning. However, placing constraints on the availability of the information made a big difference to the age at which children demonstrated such an effect. Thus, real world cases of cumulative culture may require advanced cognitive resources which allow us to access, process, and store the information made available by others.

While examining the development of the capacity for cumulative culture is interesting, such an ability relies on more fundamental capacities for social information use or transmission. To investigate the sociocognitive capacities that might be necessary for distinctively human cumulative culture we need to look more closely at capacities for social transmission. For example, contextual differences between populations in experimental studies have shown that cumulative culture in adults is influenced by factors such as teaching (Caldwell et al., 2018) and the number of available social models (Muthukrishna et al., 2014).

In this thesis I will focus on looking at the more fundamental question of whether children are really aware of the fact that they are using social information, and their understanding of what makes that social information valuable or useful. The characteristic improvement over generations observed in adults' cumulative culture – and the fact that this appears to occur regardless of the form of the available information (Caldwell & Millen, 2009) – may even be explained by the learners' understanding of the value of the social information available from the previous generation. Therefore, if children have a limited, or lack of, understanding of this value they may be less likely to use the available social information appropriately. Thus, the degree of improvement across generations may be diminished relative to that of adults. Such an interpretation would also help to explain the absence of evidence of cumulative culture in non-humans inasmuch as recognising the value of social information is likely beyond their capabilities. That is, if recognising the value or relevance of social information is necessary for human-like cumulative culture then that might explain why non-humans show evidence of cumulative culture only in limited, context-specific, cases (where such an understanding might not be necessary). Fundamental questions regarding the developmental trajectory of children's understanding of social information use must be tackled before we can truly attempt to answer questions about children's (and by logical extension non-humans') potential to demonstrate the kind of cumulative culture observed in adults. The following section reviews a number of sociocognitive mechanisms that have been proposed to underlie distinctively human cumulative culture.

Prerequisite Sociocognitive Mechanisms

By definition, cultural transmission necessitates using information acquired from a social source. In the current thesis, when investigating the sociocognitive mechanisms proposed to underlie distinctively human cumulative culture I have focused in particular on the sophistication of social information use. In doing so, I have been able to study the mechanisms that may be implicated in cumulative culture in a way that permits reasonable comparisons across populations. Learning from others – social learning – is widespread in humans and a variety of non-human species. However, there appear to be key differences in human adults' social learning (relative to non-humans) that could account for the distinctiveness of human cumulative culture (Boyd & Richerson, 1996; Dean et al., 2014; Heyes, 2016c). These differences are evident in human adults' capacity to appropriately seek out, attend to, and use social information across a vast array of contexts, posing a potentially significant advantage for cultural transmission (Morgan et al., 2012).

A number of sociocognitive mechanisms have been proposed as being what sets human learning apart from that of non-humans (Boyd & Richerson, 1996; Dean et al., 2014; Lewis & Laland, 2012; Tennie et al., 2009). To give an example, Dean et al. (2014) proposed that human cumulative culture is the consequence of a 'suite' or 'package' of sociocognitive abilities that enable high-fidelity information transmission. The proposed package of capacities included imitation, teaching, complex communication, and prosociality. However, this proposal does not account for the absence of evidence of cumulative culture in non-human species which have shown evidence of some of these abilities. Thus, the disparity in evidence of cumulative culture between humans and non-humans, in spite of many common sociocognitive capacities, presents an interesting conundrum. In this section I will review some of these proposed sociocognitive prerequisites and I will highlight the proposal that informs the experimental framework applied to the studies reported in chapters to follow.

The most prominent cognitive capacity that has been proposed in an attempt to explain the distinctiveness of human cumulative culture is 'imitation' or action-copying – the proclivity to copy the actions others (Tomasello et al., 1993). Humans appear to have a particular aptitude for copying the actions of others (Heyes, 2011), and the resulting high-fidelity information transmission has been claimed as being what permits cumulative culture (Tennie et al., 2009; Whiten, 2017b). Indeed, higher-fidelity cultural transmission has been found to lead to exponential growth in trait longevity, and lower rates of trait loss (Enquist et al., 2010; Lewis & Laland, 2012). However, claims that imitation is a learning mechanism that can explain cumulative culture are called into question when we consider evidence of imitation in non-humans (for a review see Huber et al., 2009) which exhibit limited evidence of cumulative culture. For example, there is evidence of imitation in chimpanzees (Whiten et al., 1996), marmosets (Voelkl & Huber, 2000), and wild birds (Aplin et al., 2015). Furthermore, evidence for cumulative culture is found in human adults even without the opportunity to imitate (Caldwell et al., 2012; Caldwell & Millen, 2009) suggesting that imitation is not essential for cumulative culture. Although, it seems likely that action copying may support cumulative culture when the behaviour required to achieve the end state is cognitively quite opaque (e.g., as proposed by Wasielewski, 2014). The evidence of sociocognitive capacities such as imitation in non-humans suggests that it is unlikely that any single capacity, such as this, is solely responsible for human cumulative culture.

An alternative proposal focuses on when and how social information is used, rather than a particular type of social information. Social learning is generally assumed to be an adaptive capacity when it is present within a population. Yet although it is beneficial and tends to be more efficient relative to individual learning (Rendell et al., 2011), theoretical models suggest that the ability to be selective and flexibly switch between social and individual learning is what increases efficiency and adaptability (Enquist et al., 2007; Henrich & McElreath, 2003; Rogers, 1988). That is, relying on *only* social learning does not present the opportunities for innovations to occur within a population which are necessary for improvements to a cultural trait to be accumulated. Recent evidence, drawn from datasets generated by collaborative online programming competitions, indicates that selective social learning produces cumulative improvement in knowledge over time (Miu et al., 2020). The results of this study showed that those who performed best employed a mixture of copying those who were more proficient, and learning individually through adding their own innovations. In other words, copying alone cannot increase adaptability in a population, relative to individual learning. Likewise, always innovating is also likely to lead to relatively low success, and accumulation of benefits over generations would not be possible. Rather, it is the ability to flexibly switch from copying to individual learning, when copying proves unsatisfactory, and vice versa, that increases adaptability.

Another way to think about selective social learning is in terms of targeting learning towards appropriate social sources, specifically those who possess relevant knowledge or experience. To some extent this selective, or targeted, social learning has been observed in humans and non-humans in the form of social learning strategies (SLSs). These are flexible rules that direct learning towards objects, agents, or events that are most likely to provide useful information (Kendal et al., 2018; Laland, 2004). SLSs help to filter out less useful aspects of available social information, influencing selectivity in social learning. These strategies have been considered responsible for much of the adaptive social learning found in humans and non-humans (Kendal et al., 2018; Rendell et al., 2011). The selective nature of SLSs such as ‘copy older individuals’ or ‘copy the majority’ makes them generally more effective than learning indiscriminately from others or individual learning through trial and error. However, evidence of SLSs, and selective social learning more generally, in both humans and non-humans (e.g., Price et al., 2017; Wood et al., 2013b) suggests that selective social learning also cannot fully account for the differences between these populations in regard to the capacity for cumulative culture.

There is however an alternative proposal regarding strategic social learning, which, in contrast, does focus on capacities widely regarded as unique to humans. According to this proposal, the distinctiveness of human cumulative culture can be attributed to humans capacity to use explicitly metacognitive SLSs (Dunstone & Caldwell, 2018; Heyes, 2016c). These are consciously represented and reportable rules that reflect an explicit awareness and understanding of the reasoning-based learning strategies being employed. Explicitly metacognitive SLSs could therefore enable learners to recognise the potential value to themselves of information provided by others. This could be due to a reasoned understanding of why a particular person is, potentially, a good source of information (Whalen et al., 2018) or the ability to recognise shortfalls in one's own knowledge or access to information. For a learner, having an explicit awareness of the strategies they are using and a reasoned understanding of the value, or potential value, of social information would allow them to flexibly identify, select, or disregard social information across varied contexts (De Oliveira et al., 2019; Dunstone & Caldwell, 2018; Heyes, 2016c). This ability to recognise the value of social information and target learning towards appropriate sources is believed to, at least in part, facilitate human-like cumulative culture (Heyes, 2016a; Shea et al., 2014).

Explicitly metacognitive SLSs form one half of a dual-process account of social learning that suggests that there are distinct categories of SLSs, each based on specific types of decision rules (Heyes, 2016c). So, while human adults are assumed to have the capacity to use explicitly metacognitive strategies (Dunstone & Caldwell, 2018; Heyes, 2016c), the majority of the behaviours that conform to SLSs in both humans and non-humans are proposed to arise as a consequence of general-purpose associative learning processes or biologically selected biases. Indeed, these *implicit SLSs* (as I will refer to them) are proposed to be available to both humans and non-humans, while explicitly metacognitive SLSs are only available to humans. Implicit SLSs differ from explicitly metacognitive ones in that they are not driven by a causal understanding of the potential value of social information. Rather, they direct learning towards objects, agents, and events that have in the past (in either individual experience or phylogenetic history) provided useful information. Thus, in principle, any salient association can be formed, or bias can be inherited. While implicit SLSs afford adaptive flexibility within a population they do not offer the equivalent degree of flexibility afforded by explicitly metacognitive SLSs (Heyes, 2017).

Although implicit SLSs – and SLSs defined more broadly (Kendal et al., 2018) – appear to be ‘strategic’ this is not necessarily the case in the everyday sense of the word. Non-humans can, and do, use social information in strategic ways. However, it seems unlikely that they are basing these strategies on a causal understanding of the value of the information or that they are even aware of their use of such strategies. Rather, the strategies they use are likely the result of less cognitively demanding learning processes such as associations or heuristic biases (i.e., implicit SLSs) which nonetheless function to produce behaviour which is likely to be superficially similar to that produced by reasoning-based understanding. Therefore, Heyes's (2016c) proposal that non-humans do not have the capacity to use explicitly metacognitive SLSs and instead rely on implicit SLSs seems extremely plausible. If it is true, the effectiveness of the strategies that non-humans use could be limited in many of the contexts in which real world cumulative culture occurs. As such, this proposal could offer an explanation as to why there is so little evidence of cumulative culture in non-humans. However, categorisation of cases of strategic social learning as either implicit, or explicitly metacognitive, has thus far been ascribed only retrospectively, rather than on the basis of a direct test designed to differentiate them (Heyes, 2016a, 2017; Heyes & Pearce, 2015). Therefore, research is required to assess the validity of such a proposal.

Similarly to non-humans, young children’s social learning is thought to be driven by implicit SLSs (Heyes, 2017). Thus, given the assumption that adults have the capacity for explicitly metacognitive SLSs, and that this capacity is likely to be experience-dependent (Heyes, 2016c), it follows that this understanding must develop during the course of childhood. In the chapters to follow, I will use this framework of implicit and explicitly metacognitive social information use to describe the kind of reasoning-driven social information use that I think may pose challenges for young children. If young children find this kind of social information use challenging, then by logical extension, so too would non-humans. Later in this chapter, I will return to the issue of the cognitive challenges that may be associated with social information use, and how these could link with age-related advances in children's cognitive capacities.

Investigating the scope of selective social learning in children may help us to explain the origins of distinctively human cumulative culture. If I can identify when children develop the capacity to use explicitly metacognitive over implicit SLSs I may be able to use this to predict the likelihood of the capacity being available to non-humans. That is, if such an ability develops late in childhood then it may depend on cognitive mechanisms that are not available to non-humans and could help to explain the distinctiveness of human cumulative culture. In the following section I will review evidence of cognitive development that might be relevant to the development of the capacity for cumulative culture in children.

Cognitive Development

To establish which cognitive mechanisms could be key to the distinct differences between human and non-human social learning it is necessary to understand their ontogeny. Cognitive abilities expand over the course of childhood and adolescence, with some of the most dramatic advances in cognitive development occurring in early childhood. In this section I will introduce a few of the most pertinent cognitive abilities. I will outline their ontogeny, highlighting particular developmental milestones, and overall will provide evidence of the continuing advancement of cognition from infancy through adolescence.

The first comprehensive theory of cognitive development was set out by Jean Piaget (Barrouillet, 2015). Piaget's observations of the kinds of mistakes children made at different ages indicated to him that there were underlying cognitive mechanisms that facilitated understanding, and as these mechanisms developed they would result in particular changes in children's ability to reason about the world. This conclusion led him to propose that children pass through a number of observable developmental stages on their way to adult-like cognitive logical thinking and reasoning abilities. Specifically, Piaget described four universal sequential stages (sensorimotor, preoperational, concrete operational, and formal operational) which every child goes through at approximately the same age. Each of these stages represented a distinct capacity for, or way of, reasoning which children could only achieve once they had reached that stage. However, although these stages were described discontinuously there are not strict dividing lines between them, the process of development being continuous (Flavell, 1971). That is, it is an ongoing process in which infants and children develop and adjust their understanding of the world based on experience. Evidence to support Piaget's theory was initially limited, and when tests were developed to assess the feasibility of his theory much of the evidence pointed to Piaget having been too conservative in his predictions of children's cognitive capacities at each age. For example, later research using paradigms such as the violation of expectations task indicated that infants were capable of understanding object permanence earlier (around 5 months; Baillargeon et al., 1985) than Piaget had predicted through his own observations (around 8 months). Similarly, with regards to tasks of conservation (understanding that superficial changes in the appearance of a quantity does not mean that the quantity itself has changed), later studies discovered that slight changes in the set-up of the tasks (e.g., accidental transformation; McGarrigle & Donaldson, 1974) led to successful conservation in younger children (4 and 5 years) than was found in the original tasks (from 6 years). This suggested that factors such as the language used in the tasks could influence the outcome, something that had been neglected in Piaget's theorising (Barrouillet, 2015; McGarrigle & Donaldson, 1974).

Though different theories regarding the ontogeny of cognition present their own developmental timelines and proposals for underlying mechanisms, there is little dissent regarding the basic idea that cognition advances cumulatively through childhood and adolescence. Research into cognitive development has certainly moved on since Piaget set out his original theory. However, much of the work that has been carried out during the intervening period does link back to the idea that children's cognitive development is a cumulative process characterised by a cycle of change and growth. While generally considered plausible at the time, today there is much less conviction about, and emphasis on, the notion of stages of cognitive development as a whole. However, there are more specific developmental 'stages' related to particular cognitive mechanisms such as theory of mind, metacognition, and executive functions.

While cognition continually advances throughout infancy and childhood and into adolescence, there are what some might call key developmental milestones. These milestones represent what appear to be important shifts in children's cognitive abilities. For example, there is a general consensus that children's theory of mind – their understanding of others' mental states – undergoes an important shift at around 4 years old. This is observed most readily in children's ability to pass standard tests of false belief. In such tests, children are asked to predict where an agent, who is mistaken about an object's location, will look for the object. Children reliably answer this question correctly from around 4 to 5 years old, whereas younger children do not (Krachun et al., 2009; Wellman et al., 2001; Wimmer & Perner, 1983).

While the first experiment to look at theory of mind was conducted with chimpanzees (Premack & Woodruff, 1978) the challenge of assessing the attribution of mental states, both to oneself and others, was soon applied to human cognitive development (Wimmer & Perner, 1983). The premise for the earliest tests of theory of mind was that one way to establish whether an individual has the concept of belief would be to assess whether they can attribute a false belief (i.e., a belief or representation that contrasts with reality; Dennett, 1978; Spaulding, 2020). As outlined above, classic tests of false belief such as the 'unexpected transfer task' (Wimmer & Perner, 1983) and the 'unexpected contents task' (Gopnik & Astington, 1988) concur that by around 4 to 5 years old children develop the ability to attribute false beliefs to others. There are many researchers who contest the notion that children's false belief understanding develops as late as 4 years old. Indeed, there are a raft of studies conducted with infants which claim to have found evidence of theory of mind using spontaneous responses, used as indirect indices of understanding, rather than directly elicited responses (such as answers to explicit understanding-based questions). These spontaneous measures include anticipatory looking, looking time, and violation of expectation paradigms (e.g., Onishi & Baillargeon, 2005; Southgate et al., 2007). These tasks are designed to minimise task demands associated with the standard false belief tasks, which rely on verbal communication, and facilitate testing in non-verbal or pre-verbal populations. Sometimes these spontaneous measures are referred to in the literature as implicit theory of mind tasks. Tasks such as these report that false belief understanding may be present from infancy (e.g., from 7 months, Kovács et al., 2010; 15 months, Onishi & Baillargeon, 2005 though see Baillargeon et al., 2010 for a full review). Spontaneous measures have also been used in tests with great apes (Kano et al., 2019; Krupenye et al., 2016). An alternative to these tasks is the 'active helping' paradigm. This paradigm does not require verbal responses, and does not involve a direct understanding-based question, but rather than relying on measures of processing during passive observation (e.g., tracking eye gaze during false belief scenarios), it does require an active response – assisting an agent with their goal (Buttelmann et al., 2009; Priewasser et al., 2018; Warneken & Tomasello, 2006).

However, there has been a lot of dispute about whether these spontaneous measures actually capture early theory of mind understanding. Recent issues regarding failed replications suggest that these spontaneous measures, regardless of their validity, are also not robust or reliable (see Kulke et al., 2019). It is also difficult to know exactly what infants or great apes are responding to. For example, some researchers have argued that spontaneous behavioural response tasks highlight only what an individual sees as unusual, which does not necessarily require an understanding of mental states (Buttelmann et al., 2009; Perner & Ruffman, 2005). Furthermore, on the assumption that the tests may be tapping a different type of understanding from the explicit tests, the implicit nature of this understanding is unlikely to underpin the kind of theory of mind which could be unique to humans. Thus, despite the debate regarding spontaneous measures of belief understanding, the general consensus is that by around 4 to 5 years old children understand that others can hold conflicting beliefs and can attribute them accordingly. Indeed, 5-year-olds, but not great apes, have been found to successfully pass other non-verbal tests of false belief (Call & Tomasello, 1999; Krachun et al., 2009).

Theory of mind has been proposed to be linked to metarepresentation. If that is the case, then the evidence is more consistent with the theories that propose late-development than those that propose early-development. The attribution of mental states in theory of mind has been proposed to be based on a more general capacity for *metarepresentation*. Metarepresentation refers to the representation of representational states, and a mental state can be thought of as a representation of the world (i.e., a representational state; Perner, 1991). While representation of others' mental states is akin to theory of mind, representation of one's own mental states is generally referred to as *metacognition* (Carruthers, 2009). Metacognition can be broadly described as the ability to recognise and monitor knowledge of one's own cognitive processes (Flavell, 1979). This ability relates to understanding one's own states of knowledge and ignorance, but also an awareness of how knowledge is formed (Kuhn, 2000). It is difficult to truly separate theory of mind and metacognition. Indeed, there have been a number of accounts describing the relationship between these two capacities (for a review see Carruthers, 2009). The account that appears to have most traction proposes that metacognition might rely upon prior theory of mind understanding. That is, metacognition is the result of focusing our theory of mind capacities upon ourselves.

Metacognitive capacities appear to emerge early in children's development and these early capacities appear to be shared with non-humans (e.g., chimpanzees, Beran et al., 2015; rats, Foote & Crystal, 2007). However, during development metacognition becomes more explicit and effective as it becomes increasingly under control (Kuhn, 2000). Behavioural tests of implicit metacognition commonly employ 'opt out' (Carruthers & Ritchie, 2012) or 'ask for help' (Goupil et al., 2016) paradigms to assess individuals' appreciation of uncertainty in decision making. For example, Balcomb and Gerken (2008) reported that 3.5-year-old children engage in memory-monitoring, as shown by their tendency to skip trials in which they were uncertain, suggesting that children may have an implicit awareness of their own knowledge states from a very young age. However, the ability to recognise uncertainty is not akin to full metacognitive understanding – a reflective evaluation of one's own knowledge (S. R. Beck et al., 2012). Only from around 6 years old have children been found to accurately report what and how much they know. Younger children repeatedly overestimate how much they know in both verbal (Wimmer et al., 1988) and behavioural (Kloo et al., 2017; Rohwer et al., 2012) partial exposure tasks (participants are exposed to a range of objects, but then cannot see which of the objects is being put inside a box). Similarly, research into the development of children's response to uncertainty has indicated that 5- to 6-year-old's ability to evaluate their own knowledge could be influenced by how easily they can imagine multiple possible outcomes (S. R. Beck et al., 2011). For example, if the identity of a hidden object is known, children might find it easier to imagine a particular outcome, and therefore may only acknowledge (guess) one of the multiple possible outcomes. However, children's explicit judgements about their confidence in such guesses repeatedly reveals unwarranted overconfidence (S. R. Beck et al., 2012) suggesting that children find it difficult to evaluate what they know. It seems unlikely then that the early forms of metacognition, shared with non-humans, are sufficient to support reasoned understanding of one's own learning needs as is demonstrated in human adults. That is, although they may demonstrate some aversion to risk (e.g., in opt out tasks) or recognition of knowledge or ignorance, they do so apparently without explicit awareness of those states. This early metacognition appears to precede the ability to explicitly evaluate knowledge states in young children (Robinson et al., 2008; Sodian et al., 2006). Thus, while responses may give the appearance of an understanding of states of knowledge, the inability to report these states suggests otherwise. Rather, these early metacognitive abilities may be simply the foundation for the more explicit metacognition that develops later. This explicit metacognition has been proposed to be unique to humans (Metcalfe, 2015) and potentially implicated in distinctively human cumulative culture (Dunstone & Caldwell, 2018).

Executive functions play an important role in learning and problem solving. Studying their developmental progression in childhood offers a key insight into the advancement of cognitive development more generally. Also referred to as cognitive control, executive functions refer to a group of core cognitive processes that are necessary for self-regulatory behaviour such as directing attention when it would not be appropriate to react automatically (e.g., goal-directed behaviours, Diamond, 2013; Miller & Cohen, 2001). The three core executive functions are generally agreed to include: working memory, inhibitory control, and cognitive flexibility (Miyake et al., 2000). Unlike some other cognitive abilities, the protracted development of executive functions extends through adolescence (Anderson, 2002; Best & Miller, 2010). Indeed, there is extensive evidence of continued advancement of executive functions in late childhood and adolescence (Brydges et al., 2014; Davidson et al., 2006; Xu et al., 2013). Although the core executive function components rely on related and highly correlated cognitive processes, each develops according to its own trajectory (Anderson, 2002; Diamond, 2006).

Inhibitory control, or inhibition, can be described as the ability to deliberately control attention and inhibit automatic or prepotent responses when necessary, such that reactions and outcomes are intentional (Miyake et al., 2000). Evidence suggests that inhibitory control develops dramatically in infancy and early childhood (Carlson & Moses, 2001). Tasks such as the day-night task (Gerstadt et al., 1994) work by assessing children’s ability to inhibit a prepotent response (e.g., saying “day” when viewing a sun) and responding instead with an alternative response (e.g., saying “night” when viewing a sun). In such tasks by around 6 years old children are reliably able to inhibit the prepotent response, while younger children (3 to 4 years) make more errors and take longer to respond. This pattern of performance, also found in other similar tasks (Carlson & Moses, 2001), shows that there is a gradual improvement in inhibitory control between 3 and 6 years. However, this component continues to develop through adolescence as evidenced by tasks such as the Stroop test, in which participants’ performance continues to improve until age 21 (Best & Miller, 2010). Working memory refers to the ability to store and manipulate information over brief periods of time when it is no longer externally available (Baddeley & Hitch, 1974; Best & Miller, 2010; Diamond, 2013). Working memory has been found to improve in early childhood (Gathercole et al., 2004). Gathercole et al. (2004) found that the developmental trajectories of simple and complex working memory tasks were similar, describing a linear increase from 4 to 14 years old which flattens out between 14 and 15 years old. This suggests that working memory capacities develop progressively through childhood. Finally, cognitive flexibility, sometimes referred to as set shifting, involves thinking about something in multiple ways, for example, task switching or considering someone else’s perspective on a situation (Miyake et al., 2000; Zelazo, 2015). Though successful shifting has been found in limited contexts in children aged 3 and 4 years old (Hughes, 1998b), the most rapid development of cognitive flexibility appears to occur between 7 and 9 years old (Anderson, 2002). Specifically, 7-year-olds struggle with switching in multi-dimensional switching tasks, however this ability improves considerably by 9 years old, and continues to improve into adolescence (Anderson et al., 2000).

A common test of executive functions in young children is the Dimensional Change Card Sort task (DCCS; Zelazo, 2006). This task serves as a relatively comprehensive measure of the core executive functions in early childhood, requiring cognitive flexibility, working memory, and inhibitory control (Zelazo, 2015). The DCCS task involves children being shown two target cards (e.g., a blue rabbit and a red boat), they are then asked to sort a series of cards (that differ on the dimensions of colour and shape, e.g., blue boats and red rabbits) according to one of those dimensions. Then in a post-switch phase, children are asked to sort the same set of cards according to the opposite dimension. During this task, cognitive flexibility is required to make the shift between rules, working memory is required to keep in mind and apply the current rule, and inhibitory control is needed to prevent the perseveration of the previous rule. Children aged 3 years tend to fail this task by continuing to sort the cards according to the first dimension even in cases when they can report the current rule to the experimenter. By around 5 years old children can switch between dimensions. However, it is important to note that this does not mean that executive functions are fully developed by this age. For example, not until around 7 to 9 years can children switch flexibly on a trial-by-trial basis as tends to be required in adult tests of cognitive flexibility (Davidson et al., 2006).

When considered broadly, executive functions develop gradually. However, individual components advance through ‘jumps’ or ‘step-changes’ in capacity. That we know the approximate timing of these more specific advances means that we can assess whether changes in other cognitive mechanisms are related to, correlated with, or dependent on such executive function abilities. For instance, we know that the three core executive functions provide the basis for higher-order executive functions such as reasoning, problem solving, and planning (Diamond, 2013). So, taking reasoning as an example, working memory is crucial for the ability to make connections between apparently independent components. Working memory also makes it possible to apply existing knowledge to current perceptual contexts such that reasoning, or decision making, can be based on understanding and past experience rather than just reacting to perceptual stimuli (Diamond, 2013).

Complex language is a hallmark of our species, and unquestionably language develops with age. While early signs of language emerge in the first year of life (e.g., babbling), adult-like language abilities do not develop until later. Children’s vocabulary develops dramatically during early childhood, estimates vary but children go from having a vocabulary of approximately 6 words at 12 months old, to having a vocabulary of 10,000 – 14,000 words at 6 years old (Bloom & Markson, 1998; Herschensohn, 2007; Saxton, 2017). While language might not be necessary for the development of cognitive capacities themselves, quite often the cognitive tests used to assess them do require a certain degree of language development, at least in that they require a reasonable level of vocabulary. There are many reports of a close relationship between false belief understanding and verbal ability (Hughes, 1998a). For example, Jenkins and

Astington (1996) found that after controlling for age, general language ability and verbal memory were found to be significant predictors of false belief understanding.

More generally, the development of any cognitive mechanism can be categorised with regard to the source of the information that is required for its development (Heyes, 2018a, 2019; Shea, 2013). That is, whether development of a particular capacity is dependent on information that has been inherited genetically (nature), information gained by direct interaction with the environment (nurture, e.g., associative learning processes or trial and error), or information gained via social interaction (culture; Heyes, 2018a, 2019). Abilities that develop in line with increased opportunities for interaction with the environment or social interaction would indicate experience-dependence (Heyes, 2018a). That is, cognitive abilities that emerge only after a learner has had the opportunity for direct interaction with the environment or the opportunity for relevant social interaction (e.g., when information is not directly available from the current environment) can be described as experience-dependent. Learning through social interactions enables learners to gain information that has potentially been constructed from the experience of many others, passed on from multiple previous generations. Therefore, information gained through social interaction may achieve learning beyond that which an individual could achieve alone.

Whether particular cognitive abilities are experience-dependent or not can be assessed by looking at their developmental trajectory. As these capacities require the opportunity to learn through experience, they are expected to develop relatively late. To give an example, learning to read can be considered as experience-dependent as although an individual may have developed requisite capacities such as language (Lundberg, 2010) and working memory (Savage et al., 2006) they typically do not learn to read without being taught by others who can read. That is, developing the ability to read depends on opportunities for relevant social interaction (Heyes, 2012, 2019) not only exposure to written words. Simply looking at written words will not enable a child to acquire the skill of reading, they also need to hear the words being read aloud in order that they can form links between sounds (phonemes) and the relevant written symbols (graphemes). We can say that reading is experience-dependent as reading ability improves with increased exposure to relevant social learning (e.g., teaching). Therefore, the ability to read emerges relatively late in development after children have had appropriate and sufficient opportunities to learn. Globally, 14% of adults aged 15 years and older (~750 million people) are considered to be illiterate (UNESCO Institute for Statistics, 2017). This rate is considerably higher in regions with low access to or poor quality of education, for example in sub-Saharan Africa and Southern Asia adult literacy rates in some countries are below 50%. These relatively high rates of illiteracy underline the fact that the ability to read (and write) is unlikely to be acquired without access to appropriate learning opportunities.

While the particular developmental trajectories of the advances described above vary both from one child to another and between cognitive mechanisms. It is clear that cognitive abilities, more broadly, continue to advance through childhood and adolescence. The repeated addition of new skills results in advances to cognitive abilities that accumulate as we get older. In this respect Piaget appears to have got it right when he proposed that children's cognitive development is an ongoing process characterised by a cycle of change and growth. In the following section I will discuss how these cognitive developments could influence social information use and how this could help to explain distinctively human cumulative culture. Furthermore, in order to highlight the advantages of studying age-related changes in social information use, I will also present a series of challenges that might relate to challenges faced in real world cases of cumulative culture.

Age-Related Changes in Social Information Use

In this section, and in this thesis more broadly, I will identify particular cognitive challenges that could be relevant for social information use in the context of real world cases of cumulative culture. Looking at the developmental trajectory of effective and appropriate use of social information in response to these particular challenges could offer crucial insights into which cognitive capacities are implicated. If these capacities emerge late in development, this may indicate that they are beyond the capabilities of non-humans and therefore could offer an explanation for the distinctiveness of human cumulative culture. That is, the challenges faced by younger children in terms of social information use are likely to be shared with non-humans. If particular contexts are identified in which cognitive developments seem to have a relevant influence on social information use, this may help to explain why cumulative culture is not often found in non-humans. Such a conclusion might then be consistent with the idea that a lot of the real world contexts in which we see cumulative culture in humans could, in practice, involve a lot of these challenges.

For instance, one challenge that might pose a problem for young children (and non-humans) is the ability to reflect on one's own knowledge (metacognition) and importantly to recognise gaps in knowledge to identify one's own learning needs. Without recognition of what information is needed, it is possible that useful and relevant information will not be exploited. On the other hand, having an understanding of what one does not know may make individuals more sensitive as to where to source useful information and what to do with it. This is particularly important in cases in which gaining relevant information through direct experience with the environment (i.e., through trial and error) would be inefficient or costly. For example, some foods require particular preparation techniques in order to remove toxins and make them safe for consumption (e.g., cycas seeds, W. Beck, 1992; manioc tubers, Henrich, 2016). These, often multistage, processes have to be learned from knowledgeable others, since relying on learning through trial and error could result in serious illness or death before a safe method of preparation is discovered. However, such information about processing methods is unlikely to be readily available if you are in a situation in which you might need it. If for instance you discovered a new plant that looked like it might be edible, but you were not sure whether this was the case, nor how it might be processed, then it would be a good idea to seek out information from someone else who has this knowledge. However, as the processing is likely performed in a different location from the food source, this is not information that you would necessarily chance upon. Therefore, you need to have an awareness of what you need to know. If you do have this understanding, you might recognise that you could look out for someone collecting the same plant and then either ask or observe them to find out whether and how it can be eaten. Thus, there are real world contexts in which appropriate social information seeking might be essential for survival.

Relatedly, consider a situation in which there are a number of people but only some of them have access to useful information – through experience or perceptual access. This might be a situation in which appropriate social information use, and consequently cumulative culture, is dependent on understanding what others know (theory of mind). That is, in order to make best use of the available information it is necessary to be able to identify which of the available others has the required knowledge. Having the ability to attribute mental states to others may offer a significant advantage for understanding who can provide useful and relevant information. Such an advantage might relate to real world contexts of cumulative culture in that identifying and then using information from an appropriate social source (such as someone who has access to relevant information) could lead to a better outcome than would have been achieved by using information provided by someone who did not have access to relevant information. Theory of mind understanding may also help overcome challenges associated with opaque information. By this I mean that being able to attribute mental states to others could make it possible to infer the outcome of others' behaviour even if the outcome itself is not directly observable. An everyday example of this might be: you are waiting for a train in a foreign country, following an announcement over the loudspeaker all the other passengers, who are on the same platform as you, move together to a different platform. As you do not speak the local language you did not understand the announcement, however, you can infer that the announcement probably conveyed that the train you are waiting for is now due to arrive at a different platform. This leads you to follow the other passengers to the new platform. The ability to form such an inference likely depends on some degree of understanding of the intentions of the other passengers and recognition that the other passengers were likely to have understood the announcement, therefore giving them access to information that you do not have access to.

Here I have provided examples that were intended to highlight situations in which having, or not having, an explicit understanding (e.g., of what others know) might make a difference to how social information is used in a particular context, in terms of setting oneself on the right or wrong track. More generally, if we see a gradual development in appropriate information use in a particular context, it might indicate that children's performance is improving due to their increasing general executive function capabilities. However, if we see that there are jumps or step-changes in performance at particular ages then this may imply that children are making use of a specific cognitive capacity that comes online at that age which then allows them to perform better. Looking for age-related changes in social information use is a useful way of identifying the potential for social transmission, and also for looking at the understanding of that transmission. However, children may demonstrate appropriate information use even when they have not in fact understood the relevance of the social information and are not aware of its potential value to themselves. As such, if one is particularly interested in the question of explicit awareness, it is important to design studies in such a way that we can distinguish true understanding and recognition of the value of social information, from use of the information using a broadly adaptive but less flexible strategy (e.g., implicit SLSs).

Depending on the developmental trajectories that are identified, it might be possible to make some claims about the cognitive developments that are necessary for distinctively human cumulative culture. It may also be possible to make some predictions regarding the potential of non-humans to show similar patterns of social information use, and thus the likelihood of their capacity to demonstrate human-like cumulative culture. In the following section I will outline the aims and hypotheses of this thesis and the specific research questions tackled within each chapter.

Thesis Goals

This chapter has summarised the current state of the literature with regard to cognitive capacities proposed to be responsible for the difference in cumulative culture between human adults and non-humans. As outlined above, such differences point to the potential value of investigating the capacity for cumulative culture in young children. Therefore, the goal of this thesis is to investigate fundamental questions regarding children's awareness of their social information use and what makes that social information valuable or useful. To investigate what is happening during development, we can look for age-related changes in children's response to social information as a means to understand the cognitive processes that might be implicated in distinctively human cumulative culture. If children's use of social information appears to differ between ages, or their performance advances at a particular age, this may indicate that children are employing different cognitive capacities. For example, the capacity to recognise the relevance or value of social information will likely result in more appropriate or flexible social information use relative to heuristic biases or simple associations. Therefore, the experimental studies in this thesis seek to investigate age-related changes in children's reasoning about and use of social information as a means to understand the cognitive capacities they might be employing. Specifically, the following three chapters will examine children's use of implicit and explicitly metacognitive SLSs (Dunstone & Caldwell, 2018; Heyes, 2016c). The studies reported in this thesis are mainly concerned with distinguishing cases in which social information use is driven by reasoned understanding, from situations which can be explained by a more implicit adaptive heuristic. That is, each study will explore the development of children's capacity to reason about and understand the value of information that has been provided by others, the potential value of information that can be acquired from others, and what makes others good sources of social information.

Chapter two begins my investigation into the development of children's capacity to use social information strategically, based on reasoned understanding about the value, to themselves, of information that has been provided by others. Understanding others' motivations and desires is likely necessary for recognising whether information available from them is valuable and relevant for achieving a specific goal. However, the vast majority of social learning paradigms are designed such that the demonstrator and the participant are motivated to reach the same goal. To address this, chapter two employs a novel paradigm to examine the developmental trajectory of children's capacity to understand others' goals relative to their own, their ability to use this knowledge to infer the outcomes of social demonstrations, and their ability to apply this understanding in order to use the social information strategically. In contrast to previous social learning studies, this paradigm compares both cases in which the goals of the demonstrator and the participant are aligned, and cases in which they are not aligned. Investigating whether children can account for others' conflicting motivations might offer insight into the SLSs they are employing and whether these change over the course of childhood.

To make best use of social information we must focus learning on those who possess relevant knowledge and experience, however social information is rarely acquired passively. Therefore, chapters three and four focus on the development of children's ability to assess *who* has the potential to be an appropriate source of information by examining whom they choose to seek information from. The aim of both of these chapters is to document the developmental trajectory of children's capacity to use reasoning-based explicitly metacognitive SLSs relative to alternative strategies for success which could be explained by simpler processes, including heuristic biases and associative learning.

Chapter three examines age-related changes in children’s capacity to seek out and use appropriate social information. Specifically, it looks at children’s recognition of the information they need to complete a particular problem and their understanding about who can provide valuable information about that problem. The methods employed in this study facilitate investigation into age-related differences in the SLSs children are employing. By looking at children’s choices of who to seek information from, it may be possible to determine whether their selections are based on an understanding of the potential value of the information the model could provide (in this case, access to relevant information), or on superficial yet salient characteristics such as age or gender.

Chapter four investigates whether children can take into account reasoning about others’ mental states when judging ‘who knows’. In particular, it looks at the ability to identify an appropriate source of information based on a causal understanding of the link between the informants’ knowledge or experience and their value as a source of information. Using a novel methodology adapted from studies of social understanding in non-human primates, this study attempts to differentiate the learning processes that underpin children’s selections of social sources. To assess children’s capacity to use explicitly metacognitive SLSs this study involves inducing a scenario in which explicit reasoning about others’ knowledge directly conflicts with an association or rule formed on the basis of more salient, but not causally relevant, cues. By incorporating a switch in the scenario, the design is intended to reveal the underlying rule or bias that participants are using.

Chapter five summarises the key findings of the empirical studies in relation to existing literature, presents a discussion of the methodological limitations, and offers suggestions for future research.

Chapter 2:

Taking Account of Others' Goals in Social Information Use: Developmental Changes in 3- to 7-Year-Old Children

---

The following chapter has been submitted for publication at the Journal of Experimental Child Psychology and is currently under revision for resubmission, the reference is given below. The chapter is presented in its submitted form. Data and analysis code are available in the OSF repository: [https://osf.io/xrwgq/?view_only=990e63c30e4d479c96868af95fdf2288].

Blakey, K.H., Atkinson, M., Rafetseder, E., Renner, E., and Caldwell, C.A. (Under Revision). Taking account of others' goals in social information use: Developmental changes in 3- to 7-year-old children.

Abstract

The majority of social learning paradigms are designed such that a demonstrator and participant have the same goal. In reality it may be relatively rare to opportunistically encounter social information provided by another who shares a goal. Rather, human adults use their understanding of others' goals, relative to their own, to strategically select and use the most relevant social information. We explored age-related changes in 3- to 7-year-old children's ability to account for others' goals relative to their own, assessing their understanding of others' goals, ability to interpret others' behaviour, and ability to apply this to their social information use. Children observed a social demonstration by a puppet whose goal aligned or did not align with their own. The demonstrator selected a capsule, peeked inside, and chose to accept or reject it, following which children made their own selection. The target capsule was dependent on the individual's goal. Children had to use their understanding of the demonstrator's goal and reaction to the capsule to infer the value of the information relative to their own goal. The results revealed a pervasive age effect on success across all aspects of the task, suggesting that taking into account others' goals in the interpretation and use of social information is not trivial. This late development is consistent with accounts proposing that such abilities may be linked to the distinctiveness of human cumulative culture.

Introduction

Experimental social learning paradigms are almost exclusively designed such that a demonstrator and participant are motivated to reach the same goal (e.g. Flynn & Whiten, 2013; Haun et al., 2012; Renner et al., 2020; van Leeuwen et al., 2018). In reality it is likely to be relatively rare to encounter social information from another individual who shares our immediate goal, at the precise moment when we are simultaneously in a position to put this to use. Effective social learning may therefore require an ability to interpret others' behaviour in relation to their goals in order to identify information that may be relevant to us, or the contexts in which it might be applied. Furthermore, it may be the case that, even when a potential learner's goal conflicts with that of a potential demonstrator, their choices and preferences may still provide valuable information that could guide the learner's responses, as long as it is possible to take account of the conflicting motivations.

Therefore, it is necessary to consider whether the ability to understand and represent others' goals is influenced by the degree to which they align with one's own. Reasoned understanding of others' goals (or intentions) is thought to be an important precursor for recognising why another's behaviour is relevant. In turn, this is likely necessary for interpreting the outcome of others' behaviour, and crucially, for using social information strategically to achieve a specific goal (Hawthorne-Madell & Goodman, 2019). Human adults regularly make use of their understanding of others' intentions, relative to their own, to strategically select and use the most relevant social information (Caldwell, 2018; Vélez & Gweon, 2019). However, there is little evidence of such capacities in children or non-human animals (henceforth animals). Thus, investigating the developmental trajectory of these abilities, and the cognitive challenges they pose to children, could advance our understanding of the differences between human and animal social learning. Evidence of relatively late development of particular ways of using social information in children, that we have little evidence of in animals, could indicate that they influence distinctively human cumulative culture (Caldwell et al., 2020; Dunstone & Caldwell, 2018; Heyes, 2016c). The lack, limited scope, or inflexibility of such capacities may have a restrictive influence on cultural transmission in animals, precluding the potential for human-like cumulative culture (Caldwell, 2018). Cumulative culture refers to the accumulation of beneficial modifications to cultural traits over successive generations of learners resulting in increased functionality or efficiency (Dean et al., 2014; Mesoudi & Thornton, 2018; Tennie et al., 2009). The aim of the current study is to investigate age-related changes in children's ability to use their understanding of others' goals to interpret the outcome of others' behaviour, and thus identify how to use the social information available to achieve their own goal. Crucially, to assess whether children's abilities to interpret and use social information are influenced by their ability to represent others' goals relative to their own, we compare instances in which the goals of the demonstrator and the participant align (i.e., same goals) and instances in which they do not (i.e., different goals).

There is strong evidence of selective social learning in both very young children (Haun et al., 2013; Wood et al., 2016) and many animal species (Horner et al., 2010; Laland, 2004). Referred to as social learning strategies (SLSs) this selectivity reflects the influence of heuristic biases on individuals' social information use, directing learning towards objects, agents, and events that are most likely to provide useful information (Kendal et al., 2018; Rendell et al., 2011). Though broadly adaptive (Heyes, 2017), even this seemingly strategic social learning does not account for the distinct differences between humans and animals in regard to the capacity for cumulative culture.

There appear to be key differences in human adults' social learning (relative to animals) that could account for the distinctiveness of human cumulative culture. These differences manifest in human adults' capacity to seek out, attend to, and use social information across a vast array of contexts, posing a potentially significant advantage for cultural transmission (Morgan et al., 2012). One proposal suggests that the distinctiveness of human cumulative culture could be attributed to the selective use of explicitly metacognitive SLSs (Dunstone & Caldwell, 2018; Heyes, 2016c). These are defined as consciously represented and reportable rules regarding an explicit awareness and understanding of the reasoning-based learning strategies being employed. Such recognition could enable individuals to recognise the potential value to themselves of information provided by others due to a reasoned understanding of why that person is, potentially, a good source of information (Whalen et al., 2018).

While adults have the capacity to use explicitly metacognitive strategies (Dunstone & Caldwell, 2018; Heyes, 2016c), much of the behaviours that conform to SLSs in both humans and animals are likely based on less cognitively demanding heuristic biases. Heyes's (2016c) dual-process account of selective social learning proposes that there are distinct categories of SLSs: *implicit SLSs* (as we will refer to them here) that are available to both humans and animals, and explicitly metacognitive SLSs that are only available to humans. Thus, the capacity to use explicitly metacognitive SLSs could be what sets human learning apart from animals.

Implicit SLSs (attributed by Heyes to be largely an outcome of general-purpose associative learning processes) differ from explicitly metacognitive ones in that they are not driven by an understanding of the potential value of social information. Experimental social learning paradigms and literature on the phenomenon of 'selective trust' report a variety of heuristic biases in children (see Heyes, 2017; Sobel & Kushnir, 2013 for reviews). Children have preferences for older over younger models (Jaswal & Neely, 2006; Wood et al., 2012), higher over lower status models (Jiménez & Mesoudi, 2019; McGuigan, 2013), and more reliable models (Zmyj et al., 2010). Arguably these preferences are the product of repeated exposure to the successes of others with particular characteristics, or previous positive outcomes from copying such others. That is, heuristic biases such as these could easily be the result of associative learning processes that result in rule-like strategies. This brings into question previous claims of 'adult-like' social information use or ambiguous claims of selective 'rule-governed' copying in young children (Heyes, 2017). The selective social information use described in children under about 5 years (see Over & Carpenter, 2012 for a review) may therefore be largely explained as a result of implicit biases.

Reasoning-based understanding of the potential value of social information is thought to be cognitively challenging and experience-dependent, emerging later in childhood (Dunstone & Caldwell, 2018; Heyes, 2016b). While individuals may appear as though they are making strategic decisions regarding social information use, we propose that younger children and animals are unlikely to be making reasoning-based decisions. Rather, they likely rely on less cognitively demanding learning processes such as associations, or heuristic biases, which nonetheless function to produce behaviour which is superficially similar to that produced by reasoning-based understanding. However, reliance on such implicit biases does not offer the equivalent degree of flexibility afforded by reasoned understanding. Indeed, implicit SLSs such as 'copy successful individuals' are only appropriate when the goals of both parties are aligned. If goals are not aligned, then copying another's behaviour or selection is unlikely to be successful for a learner even if it has been successful for another.

As outlined above, adults strategically use social information in a wide array of contexts due to their ability to recognise its potential value to themselves (i.e., explicitly metacognitive social learning). In contrast, we can reasonably assume that human infants and animals lack, or have a limited capacity for, a causal understanding of the potential value of social information. If adults can account for others' goals and experience when using social information, it follows that this ability likely develops during childhood. Yet, there have been few attempts to explore the presence or emergence of these abilities in children.

By virtue of their design, the vast majority of social learning studies involve the demonstrator(s) and the participant being motivated to reach the same goal (Haun et al., 2012; van Leeuwen et al., 2018; Want & Harris, 2001; Wood et al., 2012), even if the method of reaching the goal varies (Dean et al., 2012; Evans et al., 2018; Wood et al., 2013a). Therefore, we argue that the questions being asked in such studies are not wholly reflective of the challenges associated with social learning in the real world, as they offer little opportunity to make strategic decisions regarding whether and how to use the available social information. This failure to address scenarios in which participants are required to reason about the value, to themselves, of information provided by others, necessarily omits questions regarding flexible learning. This is problematic as models suggest that it is the ability to be selective, and flexibly switch between social and individual learning, that increases adaptability within a population (Enquist et al., 2007; Henrich & McElreath, 2003; Rogers, 1988), not only the inclusion of social learning (despite social learning generally being viewed as adaptive). Flexible social learning requires individuals to recognise situations in which social learning is most beneficial, and when it is best to learn on one's own. Yet to make such decisions, assuming these involve explicit reasoning, it is likely necessary to understand whether the available information is valuable and relevant for achieving a specific goal. We expect that this almost certainly requires some degree of representation of others' goals, as well as the capacity to interpret the outcomes of others' behaviour in relation to those goals. Thus, we propose that the ability to account for others' goals could be key to engaging in the kind of reasoning-based strategic social information use thought to be implicated in distinctively human cumulative culture. Below we discuss the developmental trajectory of sociocognitive and metacognitive abilities thought to be necessary to engage in such strategic social information use.

The development of goal understanding is likely underpinned by an understanding of the potential for differential beliefs and desires. Children's explicit understanding of conflicting mental states appears to develop between 3 and 4 years old. A study by Flavell et al. (1992) exploring young children's understanding of conflicting belief states between individuals, reported that while 4- and 5-year-olds indicated an understanding that others' beliefs could be different from their own, 3-year-olds attributed their own belief state to others. Studies of belief understanding centre on children's understanding that others can hold false beliefs (i.e., a belief or representation that contrasts with reality). Classic tests of false belief such as the 'unexpected transfer task' (Wimmer & Perner, 1983) and the 'unexpected contents task' (Gopnik & Astington, 1988) concur that by around 4 years children develop the ability to attribute false beliefs to others. However, these *explicit* tests rely on verbal communication, precluding young children and animals from completing them. Implicit tests using spontaneous behavioural measures (see Baillargeon et al., 2010 for a review) suggest that false belief understanding may be present from the second year of life and in great apes (Krupenye et al., 2016; see Krupenye & Call, 2019 for a review). However, the implicit nature of this understanding is unlikely to underpin the kind of explicitly reasoned social information use we are proposing could be unique to humans. Despite the debate regarding implicit measures of belief understanding, the general consensus is that by around 4 years children understand that others can hold conflicting beliefs. However, the development of the capacity to understand others' desires is more complicated.

The ability to recognise and represent others' goals relative to one's own has primarily been investigated with regards to understanding desires (Moses et al., 2000; Rakoczy et al., 2007; Repacholi & Gopnik, 1997; Wellman & Woolley, 1990). Repacholi and Gopnik (1997) found that 18-month-old infants were able to correctly infer others' desires for a particular type of food even when their own desire was different, in contrast to 14-month-old infants who pervasively attributed their own desire to others. However, more recent studies have been unable to replicate such an early development of this capacity (see Ruffman et al., 2017), instead finding that understanding of others' desires develops close to 3 years of age as was previously suggested by Wellman and Woolley (1990). These studies examine desire understanding in relation to emotional responses, so it is important to consider whether the developmental trajectory is the same for other categories of desires or goals (i.e., abstract or task-specific goals). There is evidence that suggests that children struggle to recognise others' desires when they conflict with their own until about 5 years (Moore et al., 1995; Rostad & Pexman, 2014).

Another way to examine children's understanding of task-specific goals is to investigate differences in understanding of others' diverse desires in cooperative compared to competitive contexts. Cooperation appears to closely align with children's ability to understand others' mental states (Priewasser et al., 2013). However, children under 5 years struggle to report another player's desire in a competitive game, instead reporting their own desire as that of the other players (Moore et al., 1995). Thus, recognising the potential for incompatible desires appears to remain cognitively challenging for longer during development in competitive contexts. Direct comparison of cooperative and competitive contexts found that 4-year-olds still have difficulty understanding others' diverse desires, and while cooperative contexts aid desire reasoning, competitive contexts appear to pose greater challenges (Jin et al., 2017). Competitive contexts may be particularly challenging due to the necessity to explicitly reason about another's conflicting desire, and how that can be used to benefit oneself. Cooperative contexts may not require the same level of explicit reasoning.

Social learning situations in which the reward value of others' behaviour is transparent do not necessitate an understanding of others' goals to achieve one's own goal. By contrast, understanding whether another's goal is aligned with one's own is particularly important in situations in which the reward value is opaque, as there are only others' reactions and choices to go on. However, an ability to understand others' goals can help you to interpret these cues. Understanding others' goals is not sufficient for interpreting an opaque outcome but is necessary to appreciate the link between experience and knowledge. That is, a learner will be required to make the inference that a demonstrator has privileged access to the information that allows them to make an informed choice. This is particularly relevant for the experimental procedure used in the present study.

Though children show evidence of understanding concepts of their own knowledge and ignorance from around 3 years (Pratt & Bryant, 1990; Rohwer et al., 2012), it is not until around 6 years that they consider the role of access to information in knowledge formation (Kloo & Rohwer, 2012). Children's appreciation of the causal influence of visual perceptual access on knowledge formation is likely to be learned, gradually developing until around 5 years (Robinson et al., 2008; Sodian et al., 2006). Prior to this children appear to rely on a 'feeling of knowing' when evaluating their own ignorance (Kloo & Rohwer, 2012), and if required to identify others' epistemic states they likely rely upon less cognitively demanding processes such as a previously formed associative rule or bias. Use of such rules or biases may give the appearance of understanding that someone is knowledgeable due to their access to information when in reality information access has not been considered, and instead they have relied on a more salient, less cognitively demanding cue (e.g., others' success).

The ability to attribute epistemic states to others based on their reaction to an opaque outcome likely relies on a sound representation of the other's goal (i.e., their desired outcome) and an understanding of whether the behaviour was knowledge based or not (i.e., whether they made an informed choice). Therefore, younger children may struggle to correctly identify the outcome of others' behaviour, despite being able to appropriately represent others' goals, due to an inability to recognise the link between information access and privileged knowledge. Even if children have developed the ability to understand and represent others' goals and use this understanding to interpret the outcome of their behaviour, they may not have the capacity to recognise the potential value of such information to themselves. To appropriately apply this understanding to their social information use, children have to recognise the value of the information relative to their own goal, an important aspect of which is determining whether the other's goal is aligned with one's own. Correctly interpreting the outcome of another's behaviour, but not appropriately applying that understanding to inform information use, may indicate reliance on implicit SLSs rather than explicitly metacognitive ones. In such cases, if learners are more likely to base decisions on salient cues such as a positive reaction to an outcome, this may appear as reasoned when their goals align with others but will result in lower success when goals are not aligned.

The aim of the current study was to investigate age-related changes in children's ability to strategically use social information by taking into account others' goals relative to their own. Specifically, we were interested in whether children's abilities to interpret and use social information were influenced by the degree to which others' goals aligned with their own. Using a simplified social learning scenario, we compared the social information use of children assigned to two different conditions: in one condition the goals of the demonstrator and the participant were aligned (Same Goals condition), while in the other their goals were not aligned (Different Goals condition). Task goals were related to retrieving a target capsule from one of two locations, each containing a different type of capsule; the target capsule was dependent on the individual's goal. Participants observed a demonstrator select and peek inside a capsule and then accept (successful outcome) or reject it (unsuccessful outcome). Both outcomes are potentially equally informative to someone who has represented others' goals relative to their own and recognised the value of the information. After observing the demonstrator's reaction, participants had to make a binary choice between *copying* the demonstrator's selection by selecting a capsule from the same location and *shifting* away from the demonstrator's selection to select a capsule from the alternative location. The appropriate response was dependent on the demonstrator's goal and the outcome of the demonstration. If the demonstrator had the same goal as the participant, the appropriate response was to copy successful selections and shift following unsuccessful selections. If their goals were different, the appropriate response was to shift after successful selections and copy unsuccessful selections. As such, task success was dependent on participants' ability to represent others' goals relative to their own and interpret the demonstration outcome to inform their own selection.

In addition to appropriate information use, we examined participants' understanding of the task goals (their own and the demonstrator's), and their ability to interpret the outcome of the demonstration. To use the information strategically children had to recognise the value of the information that had been provided in the demonstration in respect to their own goal while also taking into account the goal of the demonstrator. Therefore, even if children understood the task goals and correctly inferred the outcome of the demonstration, this would not necessarily mean that they would be able to use this information strategically.

We predicted that goal understanding, interpreting the demonstration outcome, and thus appropriate information use would improve with age because of children's developing understanding of diverse desires and conflicting motivations. Any age-related increases in children's ability to use social information strategically in this context could suggest that this ability is dependent on advanced cognitive capacities that develop relatively late. This would be consistent with proposals which suggest that explicitly reasoned SLSs may account, at least in part, for distinctively human cumulative culture. Furthermore, we predicted that performance would be greater in the Same Goals condition compared to the Different Goals condition. The expected difference between the conditions could indicate that children find it more challenging to use social information from another whose goal does not align with their own, compared to another whose goal matches their own. More specifically, differences between the conditions may offer an insight into whether children are using implicit or explicitly metacognitive SLSs. For example, low performance in the Different Goals condition compared to the Same Goals condition might indicate that children have not recognised the relevance of the social information and are relying on implicit SLSs.

Method

Participants

The final participant sample comprised 190 children aged three to seven years (97 females; Mean age = 65.7 months, *SD* = 16.6 months). Participants were unsystematically allocated to one of two experimental conditions (Same Goals condition, *n* = 97; Different Goals condition, *n* = 93) subject to the constraint that the number of children in each age group was balanced across conditions. In the final sample there were between 18 and 20 participants per age group in each condition. The nationality of the sample, as identified by parents/legal guardians, was predominantly British. Participants were recruited from nurseries, primary schools, and public attractions in Scotland. An additional 15 children were tested but later excluded from analyses due to researcher error (*n* = 3), missing data (*n* = 8), task interference (*n* = 2), or confidential disclosures by the child's guardian which cast doubt on whether that child's performance could be considered as representative of the intended recruitment population (*n* = 2). This study was granted ethical approval by the University of Stirling General University Ethics Panel (GUEP338), and informed written consent was provided by each participant's parents/legal guardians prior to data collection.

Materials

Task apparatus is presented in Figure 1. Two hand puppets, *'Rabbit'* (A.S. Puppets) and *'Bird'* (The Puppet Company Ltd.), served as demonstrators. Two round mesh buckets (342 x 235mm) each contained 23 carrot shaped plastic capsules (142 x 31mm; Figure 1A/B). Each carrot shaped capsule could be opened to reveal hidden contents (Figure 1C). The capsules inside one of the buckets contained a roll of orange felt, while the capsules inside the other bucket contained a pink knitted *'worm'*. Each bucket was fitted with an elasticated fabric lid with an elasticated hole in the centre (Figures 1A and 1B). Within a trial both buckets were fitted with lids of the same colour (e.g., both red), between trials the colour of the lids was different (red, purple, and blue). Six grey plastic tokens (23mm diameter) and a wooden box (580 x 580 x 350mm) were employed as a 'pay to play' system for the puppet demonstrators and child participants. The demonstrators stored target capsules in a wicker basket, participants stored target capsules in a plastic frying pan.

Figure 1

Illustration of the Experimental Set-Up

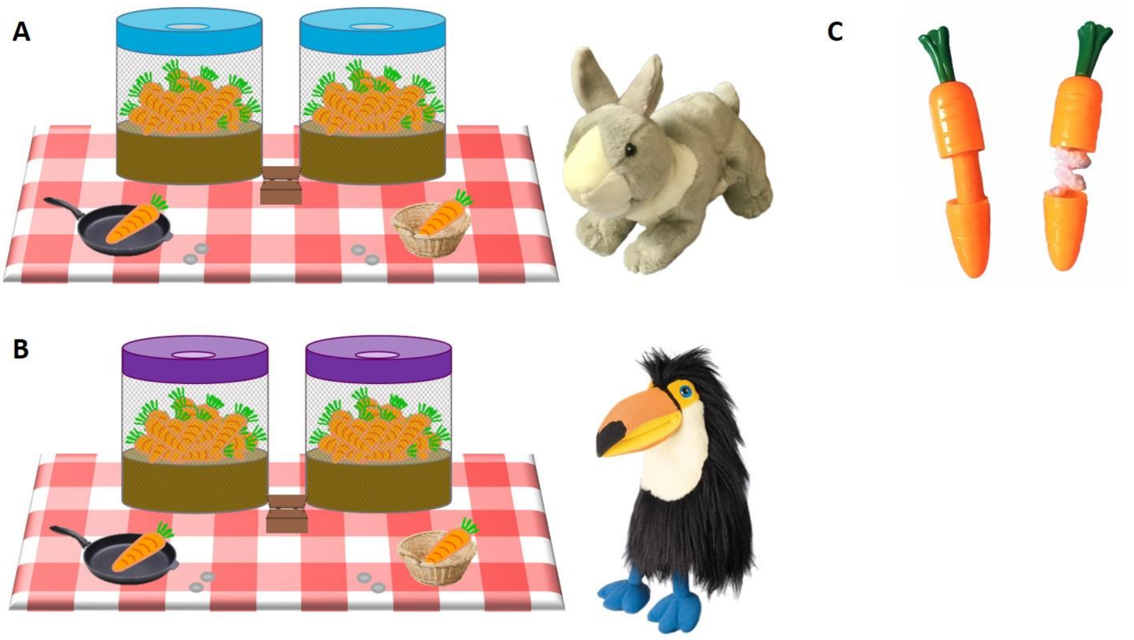

Note. A: example set-up during a trial in the Same Goals condition with Rabbit as demonstrator; B: example set-up during a trial in the Different Goals condition with Bird as demonstrator; C: example of the carrot capsules opened to reveal the two different content types, the capsules inside one bucket contained orange felt and the capsules inside the other bucket contained a pink knitted 'worm'.

Design

Participants took part individually in a single session for which they received a sticker. Two experimenters were present: experimenter one (E1) provided instructions and asked all questions (see Appendix A for a verbal script), experimenter two (E2) controlled the demonstrations, using their spare hand to assist the puppets' actions, and live coded all responses. In some sessions, a familiar adult was present but was asked not to interact with the participant.

The task comprised a series of three trials completed in one of two conditions. In the Same Goals condition the participant and demonstrator (Rabbit puppet) shared the same task goal: both searched for capsules containing orange felt. By contrast, in the

Different Goals condition the participant and demonstrator (Bird puppet) had different task goals: the participant searched for capsules containing orange felt, but the demonstrator searched for capsules containing a worm. Across both conditions the format of all trials was the same. Each trial included a social demonstration in which the demonstrator selected a capsule from one of the two buckets and *'peeked'* inside, without letting the child see the contents. Demonstrations were either successful (demonstrator selected, and accepted, a target capsule) or unsuccessful (demonstrator selected, and rejected, a non-target capsule), the target capsule was dependent on the condition. Each participant experienced a combination of successful and unsuccessful demonstrations. The combination and order of the demonstrations were randomly assigned subject to the constraint that each participant observed at least one successful and one unsuccessful demonstration across the three trials. The demonstrator accepted the capsule by placing it in their basket or rejected it by giving it to E1. Rejected capsules were given to E1 rather than being returned to the buckets to remove the risk of participants returning capsules to the wrong bucket causing potential confusion in following trials, and to avoid a noticeable difference in the quantity of capsules in each bucket, particularly in the Same Goals condition. Between each trial both buckets were removed from the table, the lids were replaced with lids of a different colour, and the buckets were returned to the table for the next trial. The sides that the buckets (distinguished by the contents of the capsules) were presented on was randomly allocated in each trial. Following completion of the third trial participants were asked to provide explicit verbal reasoning for their selections: *"How were you deciding which bucket to choose?"*.

Procedure

Participants were told that they were going to play a game with Rabbit (Same Goals condition) or Bird (Different Goals condition). Between conditions, the procedure varied only to accommodate the demonstrators' alternative goals. E1 presented the two buckets of carrot shaped capsules and explained that inside one of the buckets the carrots were orange inside and that in the other bucket the carrots had a worm inside; examples of each capsule type were presented alongside the explanation. Each trial began with the participants being informed (trial one), or reminded (trials two and three), of their own goal, *"You are looking for the carrots that have orange inside"*, and the demonstrator's goal: Same Goals condition, *"Rabbit is looking for carrots with orange inside"*; Different Goals condition, *"Bird is looking for carrots with worms inside"*. Participants were asked four goal understanding questions about their own and the demonstrator's goals (i.e., identifying contents of the target and non-target capsules). If participants did not initially respond, they were prompted with the two possible answers by E1, *"orange"* or *"worms"*. Due to a lack of verbal responses in 3-year-olds it was necessary to adapt how they were required to respond; therefore, E1 asked participants to point and held out examples of the capsule contents. Incorrect responses were not corrected. Children were considered as having goal understanding if they answered all four questions correctly.

Participants and demonstrators were each provided with three tokens. Before making a selection, each had to pay a token. The demonstrator always went first to provide the demonstration. Participants observed as the demonstrator paid a token and indicated a bucket for E2 to select a capsule from (only E2 was aware of the contents of each bucket prior to selection). Once a capsule had been selected, the demonstrator opened it slightly and *'peeked'* inside, ensuring that the contents were not visible to participants. The demonstrator then accepted or rejected the selected capsule: in the Same Goals condition Rabbit accepted carrots that were orange inside, and rejected carrots that contained worms, and in the Different Goals condition Bird accepted carrots that contained worms and rejected carrots that were orange inside.

Following the demonstration, participants were encouraged to take their turn by paying a token and selecting a capsule from one of the buckets. Once they had selected a capsule they were instructed to *"peek"* inside and choose whether they wanted to keep it in their pan or not keep it and give it to E1 (children chose to keep non-target capsules in 11.6% of trials). To determine whether children had correctly inferred the outcome of the demonstration in a given trial, they were asked a demonstration outcome understanding question: *"What was inside the carrot that Rabbit/Bird found?"* At the end of each trial participants were asked three memory questions, requiring them to recall which buckets they and the demonstrator selected capsules from, and the contents of their selected capsule.

Statistical Analysis

Analyses were performed using R (R Core Team, 2020), with generalised linear mixed effects analyses (GLMMs) performed using lme4 (Bates et al., 2014) with logit regression. P-values < .05 were accepted as statistically significant. The binary dependent variables in the analyses were: goal understanding, demonstration outcome understanding, and use of an appropriate response. Where specified as fixed effects the following variables were sum coded: condition (Same Goals as −1, Different Goals as 1) and demonstration outcome (unsuccessful as −1, successful as 1). Age was centred and scaled to measure thousands of days. The random effects structure for each model aimed to include by-participant random slopes for all fixed effects and keep random effects structures 'maximal' where possible (following Barr et al., 2013). Where the 'maximal' model resulted in non-convergent or singular fit models, random slopes were removed, followed by random intercepts where necessary, until a convergent, non-singular model was obtained. Post hoc analyses were carried out using estimated marginal means using the *emmeans* package (Lenth et al., 2019). Post hoc results are given on the log odds ratio scale.

Results

The key aim of this study was to determine whether children were able to take others' goals into account when interpreting and using social information to achieve their own goal. We were particularly interested in whether children would find it more challenging to use information provided by a demonstrator whose goal was different to their own.

Understanding

Goal Understanding

First, we looked at children's understanding of their own and the demonstrators' goals. Children were considered to have understood the task goals in 74% of the 570 trials (answered all four goal understanding questions correctly within a trial). We compared children's goal understanding with regards to the condition to which they were assigned (Same Goals or Different Goals) and age. A GLMM for goal understanding in each trial was built with fixed effects of condition and age, the interaction between these variables, and a random intercept of participant ID. A significant main effect of age (b = 6.39, SE = 1.26, z = 5.09, p < .001) indicated that children's goal understanding improved with age. A one-way ANOVA revealed significant differences in goal understanding between ages (in years), F(4, 565) = 52.28, p < .001. Post hoc comparisons using Tukey HSD tests indicated that 3-year-olds' goal understanding was significantly poorer than all other ages' (p < .001), and that 6- and 7-year-olds' goal understanding was significantly greater than 4-year-olds' ($p \leq$ .003). There was a particularly striking increase in goal understanding between 3-year-old (M = .29, SD = .45, n = 108) and 4-year-old children (M = .73, SD = .45, n = 114; Figure 2A). The GLMM did not reveal evidence of an effect of condition (b = 0.28, SE = 0.45, z = 0.63, p = .529), or an interaction between condition and age (b = 1.20, SE = 1.02, z = 1.17, p = .241) suggesting that although goal understanding improved with age there were no differences relative to the Same and Different Goals manipulation.

Figure 2

Proportion of Trials in Which Children Were Considered to Have Each Type of Understanding in Each Age Group

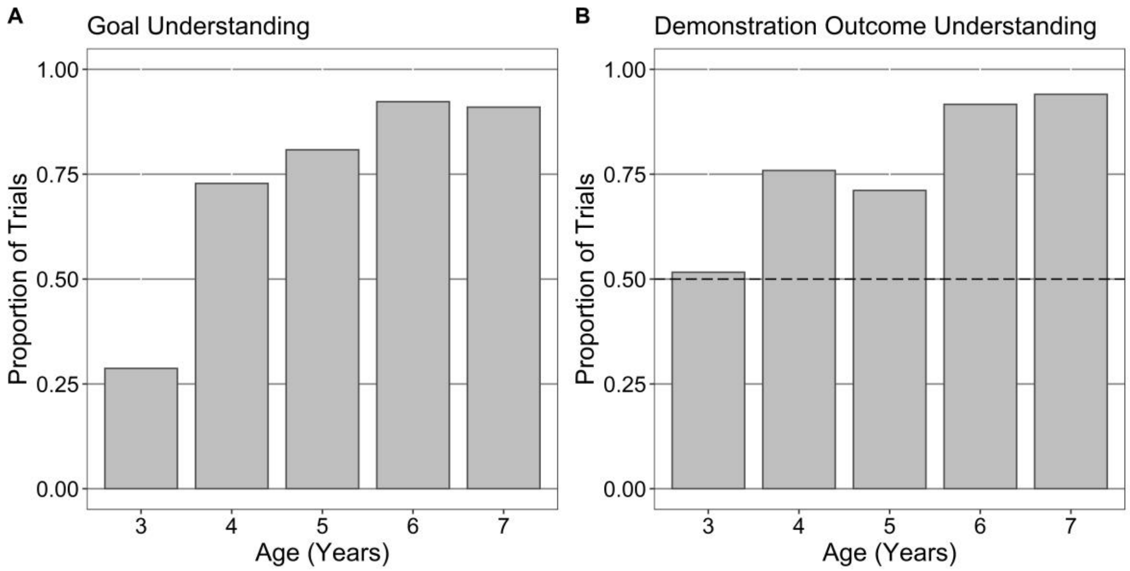

Note. A: Goal understanding (*N* = 570), children were considered to have understood the task goals in a given trial if they answered all four goal understanding questions correctly. B: Demonstration outcome understanding (including only the trials in which children demonstrated goal understanding, *N* = 420), trials in which children correctly inferred the contents of the capsule that had been selected by the demonstrator. Dashed line indicates chance.

Demonstration Outcome Understanding

Looking only at the trials in which children demonstrated goal understanding (*N* = 420), we examined children's understanding of the demonstration outcome (i.e., whether children correctly inferred the contents of the capsule that the demonstrator selected, as measured by their answer to the question *"What was in the capsule that Rabbit/Bird picked?"*). Children correctly inferred the outcome of the demonstration in 81% of trials. A GLMM was built for demonstration outcome understanding with fixed effects of condition, demonstration outcome, age, the interactions between these variables, and a random intercept of participant ID. Results revealed that children's understanding of the demonstration outcome improved significantly with age ($b = 2.19$, $SE = 0.46$, $z = 4.72$, $p < .001$; Figure 2B). Demonstration outcome understanding was significantly greater in trials with a successful demonstration ($M = .87$, $SD = .34$, $n = 202$; $b = 0.41$, $SE = 0.17$, $z = 2.40$, $p = .016$) compared to trials with an unsuccessful one ($M = .76$, $SD = .43$, $n = 218$). There was no evidence of an effect of condition ($b = 0.02$, $SE = 0.19$, $z = 0.13$, $p = .896$), nor any interactions between age, condition, and demonstration outcome ($p \geq .213$).

We analysed children's demonstration outcome understanding with regards to children's age in years. A one-way ANOVA revealed significant differences in demonstration outcome understanding between ages $F(4, 415) = 12.4$, $p < .001$. Post hoc comparisons using Tukey HSD tests indicated that 6- and 7-year-olds' understanding was significantly greater than 3-, 4- and 5-year-olds' ($p \leq .03$), and 4-year-olds' understanding was significantly greater than 3-year-olds' ($p = .016$). Binomial tests revealed that the proportion of trials in which 4- to 7-year-olds correctly inferred the demonstration outcome was significantly greater than chance ($p < .001$), while 3-year-olds were no different to chance ($M = .52$, $SD = .51$, $n = 31$, $p > .999$).

Appropriate Information Use

Appropriate information use was measured by the proportion of trials in which children selected a target capsule. Overall, children responded appropriately in 60% of 570 trials. A GLMM was built for appropriate information use with fixed effects of condition, demonstration outcome, age, the interactions between these variables, random intercepts of goal understanding and demonstration outcome understanding, and a by-participant random slope for demonstration outcome. A significant main effect of age revealed that target capsule selections increased with age ($b = 0.54$, $SE = 0.24$, $z = 2.26$, $p = .024$; Figure 3). Although information use improved significantly with age, even 7-year-olds did not perform particularly well, selecting the target capsule in only 70% of trials. Binomial tests revealed that the older three ages selected the target capsule in significantly more trials than would be expected by chance: 5 years ($p =$

.022), 6 years (p = .002), and 7 years ($p < .001$). By contrast, younger children were not different to chance (3 years, p = .387; 4 years, p = 1.00).

Figure 3

Proportion of Trials in Which Participants in Each Age Group Selected a Target Capsule

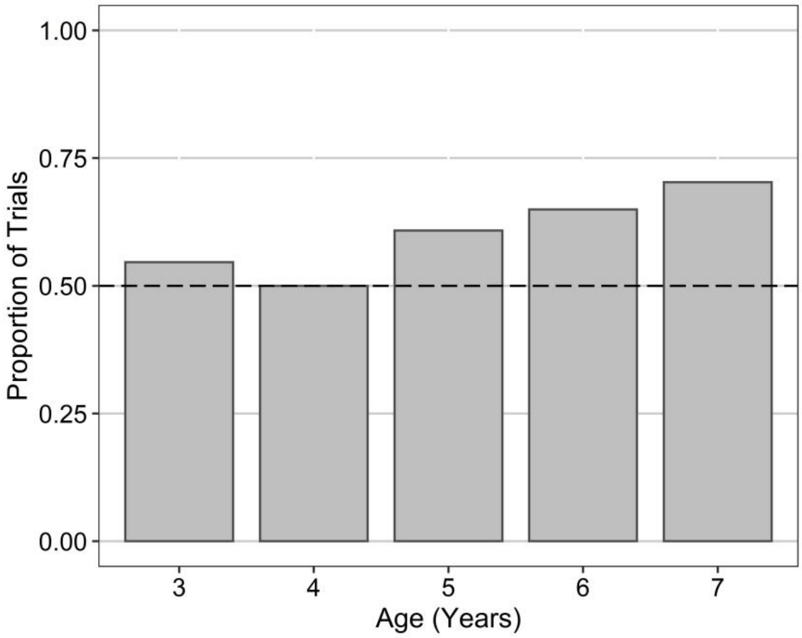

Note. Dashed line indicates chance performance.

There was no evidence of an overall effect of condition (b = −0.12, SE = 0.10, z = −1.14, p = .253) or demonstration outcome (b = −0.02, SE = 0.11, z = −0.22, p = .827) on appropriate information use. However, a significant two-way interaction between condition and demonstration outcome (b = 0.47, SE = 0.11, z = 4.44, $p < .001$) indicated that appropriate information use was dependent on both the condition and the demonstration outcome. To clarify this interaction, we performed a post hoc analysis using *emmeans*. This revealed that appropriate information use in the Same Goals condition was significantly greater following unsuccessful demonstrations (b = 0.99, SE = 0.31, z = 3.21, p = .001) and in the Different Goals condition following successful demonstrations (b = 0.89, SE = 0.29, z = 3.04, p = .002). This pattern relates to appropriate use of two possible response types: copying and shifting. Children appropriately used the copy response in only 51% of 280 trials. By contrast, children appropriately used the shift response in 69% of 290 trials (Figure 4). Shifting was the appropriate response in the Same Goals condition following unsuccessful demonstrations (71% appropriate responses in 147 trials), and in the Different Goals condition following successful demonstrations (67% appropriate responses in 143 trials).

Figure 4

Proportion of Trials in Which Children Selected the Target Capsule by Condition, Demonstration Outcome and Appropriate Response

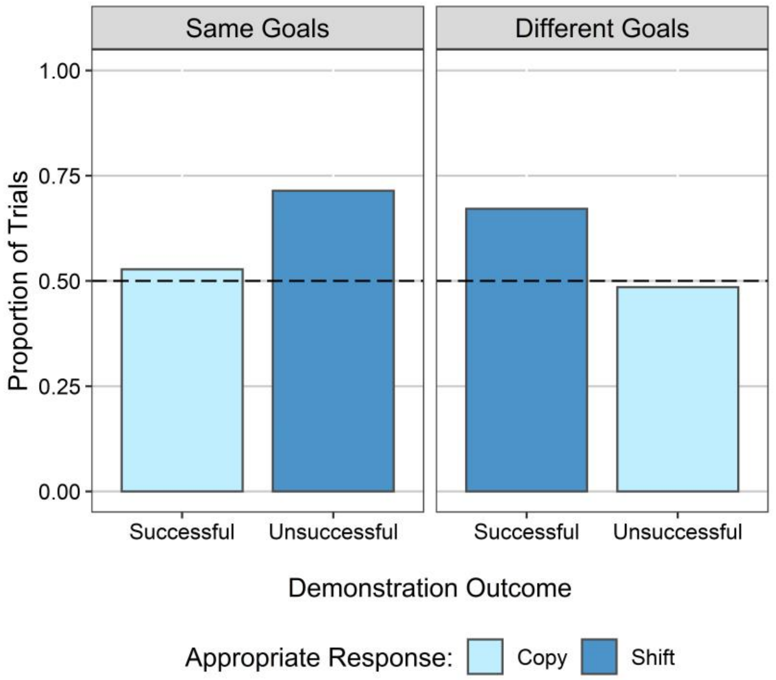

Note. Dashed line indicates chance performance.

The model also revealed a significant three-way interaction between the condition, the demonstration outcome and age (b = 0.76, SE = 0.21, z = 3.64, p < .001). To clarify this interaction, we performed a post hoc analysis using *emmeans*. This highlighted the difference between the two response types (copying and shifting) with regard to age: appropriate copying did not improve with age ($p \geq .561$), remaining close to chance, while appropriate shifting did increase with age (appropriate information use in the upper quartile was significantly greater than lower quartile; $p \leq .022$; Figure 5). Separating information use by age in years relative to the appropriate response required in each condition following each demonstration type (Figure 6) showed that 3- and 4-year-olds did not follow the same pattern of appropriate responses as 5-, 6- and 7-year-olds. It is worth highlighting that 3- and 4-year-olds' information use was no different to chance following both conditions and demonstration types, whereas 5-, 6-, and 7-year-olds information use was only greater than chance when the appropriate response was to shift. This pattern of results suggests that there may be an age-related change in the way that children approach the use of social information; in this sample this change appears to occur between 4- and 5-years-old.

For completeness, we repeated the above analysis including only the trials in which children had demonstrated goal understanding (*N* = 420), therefore we built a GLMM for appropriate information use with fixed effects of condition, demonstration outcome, age, the interactions between these variables, and a random intercept of participant ID. As expected, the results of the GLMM revealed the same effects and interactions captured in the original analysis.

Figure 5

Proportion of Trials in Which Children Selected the Target Capsule by Condition, Demonstration Outcome, Age in Months, and Appropriate Response

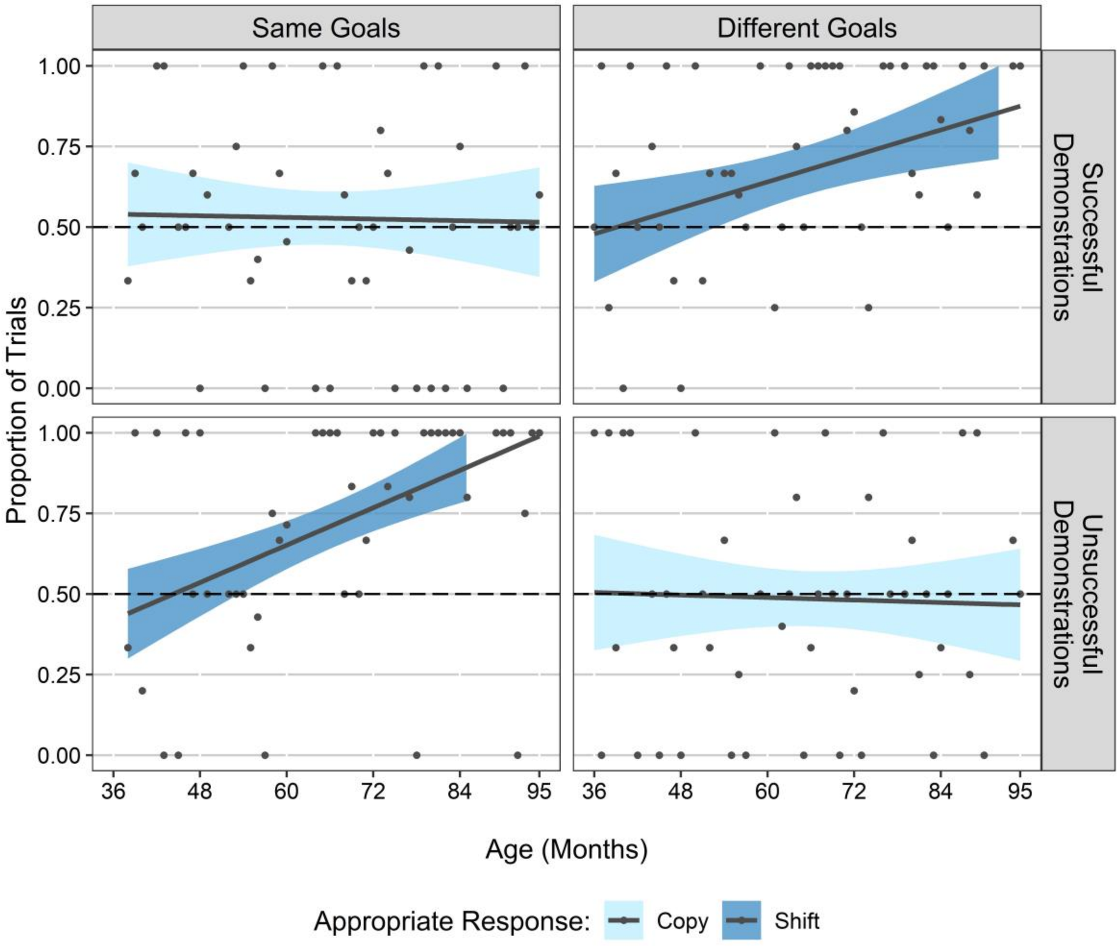

Note. Dashed line indicates chance performance.

Figure 6

Proportion of Trials in Which Children Selected the Target Capsule by Age in Years, Condition, Demonstration Outcome, and Appropriate Response

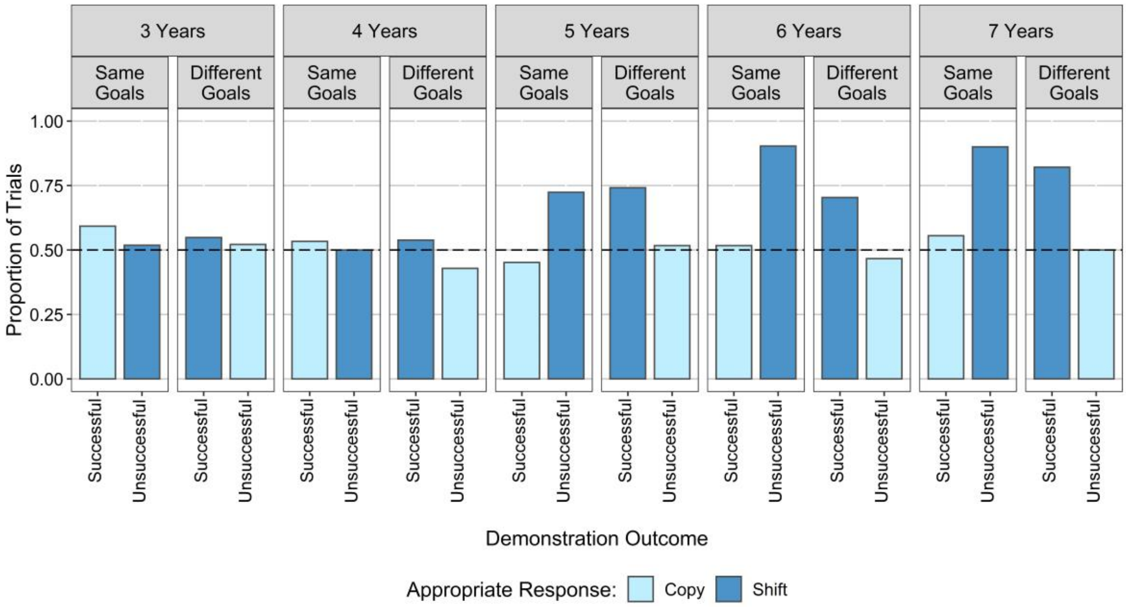

Note. The pattern of results observed in Figure 4 were only evident in 5-, 6-, and 7-year-old children, while 3- and 4-year-old children's use of both types of appropriate response was no different from chance. Dashed line indicates chance performance.

Memory Check

Children recalled the location of their own capsule selection in 93% of trials, with similar rates across all ages (84% - 97%). Children recalled the location of the demonstrator's selection in 85% of trials and showed age-related differences in recall: 5-, 6-, and 7-year-olds successfully recalled in ≥90% of trials, and 3- and 4-year-olds successfully recalled in 69% and 78% of trials, respectively. A one-way ANOVA revealed significant differences in recall between ages $F(4, 565) = 11.11$, $p < .001$. Post hoc comparisons using Tukey HSD tests indicated that 5-, 6- and 7-year-olds' recall was significantly greater than that of 3- and 4-year-olds ($p \leq .045$).

The pattern of appropriate information use for trials in which children recalled the demonstrator's selection was not different to the full sample. Binomial tests revealed that the older three age groups selected the target capsule significantly above chance (5 years, $M = .60$, $p = .036$; 6 years, $M = .65$, $p = .002$; and 7 years, $M = .70$, $p < .001$), while younger age groups were not different to chance (3 years, $M = .61$, $p = .064$; 4 years, $M = .51$, $p > .999$).

Explicit Verbal Reasoning

We examined children's responses to the explicit verbal reasoning question regarding how they were choosing which bucket to select a capsule from. Participants were asked this only after all three trials had been completed, each providing a single response. Responses were categorised into three levels: 1) non-reasoned response, including an incorrect response, an incoherent response, or no response; 2) partially reasoned response, explanation included some evidence of explicit task understanding though was incomplete; and 3) fully reasoned response, the explanation provided clear evidence of explicit task understanding.

Of the 190 participants, only seven were categorised as having provided fully reasoned responses. Older children gave more reasoned responses than younger children (Table 1). The poor rate of explicit verbal reasoning suggests that children found articulating justifications for their choices particularly challenging despite older children displaying high levels of goal understanding and demonstration outcome understanding.

Table 1

Number of Children Providing Each Category of Response to the Explicit Verbal Reasoning Question, by Age in Years

		Age (years)				
	Total	3	4	5	6	7
Non-reasoned	172	36	37	34	35	30
Partially reasoned	11	-	1	5	2	3
Fully reasoned	7	-	-	1	2	4

Adult Participants

As 7-year-olds used the appropriate response in only 70% of trials, we conducted an exploratory follow-up study using the same methods with adults (six in each condition; N = 12). Adults used the appropriate response to select the target capsule in 33 of the 36 trials (92%), and all gave fully reasoned responses to the explicit verbal reasoning question. Adults were recruited at the University of Stirling and received £3 compensation for participation.

Discussion

The methods employed within this study allowed us to examine the developmental trajectory of children's capacity to understand others' goals relative to their own, their ability to use this knowledge to infer the outcomes of social demonstrations, and their ability to apply this understanding to strategically adopt an appropriate response.

We found that the proportion of trials in which children were considered to have understood both their own and the demonstrator's task goals increased with age. A particularly large improvement in goal understanding between 3- and 4-year-olds is indicative of a developmental jump in children's capacity to recognise, represent, and report others' goals. Contrary to our predictions, whether the child's and the demonstrator's goals were aligned, or not, did not predict children's understanding of the task goals. This suggests that from around 4 years children understand others' goals even when they conflict with their own. This is consistent with evidence from the explicit false belief literature which indicates that children understand that people can hold divergent beliefs from around 4 years. Though the evidence regarding the age at which children develop an understanding of others' diverse desires is inconsistent, the lack of difference between the conditions suggests that children older than 3 years understood the demonstrators' diverse goals.

As the demonstration outcome was opaque to the participants, we examined their ability to interpret it using their understanding of the demonstrator's goal and the demonstrator's reaction to the capsule. To appropriately infer the demonstrator's reaction to the selected capsule, children needed to understand that the demonstrator had obtained perceptual access via peeking and therefore had privileged knowledge about the contents of the capsule to make an informed choice. Of those who understood the task goals, older children were found to correctly infer the outcome of the demonstration in a greater proportion of trials than younger children. Though 4- to 7-year-olds correctly inferred the outcome of the demonstrations significantly more often than would have been expected by chance, only from 6 to 7 years were they able to *reliably* infer the correct outcome. Thus, there appears to be another developmental step-change between 5 and 6 years, this time in children's capacity to interpret the content of social information.

Appropriate interpretation of the demonstration outcome likely improved with age due to two key developments. Firstly, older children were more likely to be able to use their understanding of others' goals. Secondly, older children may have been more likely to have recognised that the demonstrator had had visual access to the contents of the capsule and made the link between information access and knowledge. We can only speculate about whether this is related to a capacity to account for others' goals or related to understanding of access to information and knowledge formation. Literature suggests that children do not demonstrate understanding of how knowledge is formed with consideration of the role of access to information until around 6 years (Kloo & Rohwer, 2012). Therefore, older children's more reliable performance could be linked to the ability to reason about others' access to information. We found no effect of the Same and Different Goals manipulation on children's interpretation of the demonstration outcome, suggesting that children were not impacted by the need to account for others' conflicting goals. However, that children correctly inferred the demonstration outcome more often following successful demonstrations suggests that at least some children may have used less cognitively demanding processes such as a previously formed associative rule or bias (in this case, children may have used a positive cue from the demonstrator's choice to keep the capsule). This may give the appearance of understanding that someone is knowledgeable due to their access to information when in reality information access has not been considered. It is also possible that the younger children could have been basing their responses on a 'feeling of knowing' (Kloo & Rohwer, 2012) rather than an understanding of knowledge acquisition. That is, children under 6 years have been found to rely on the ease with which plausible information comes to mind when assessing their own epistemic states (Kloo & Rohwer, 2012). Therefore, as children had been made aware of the contents of both capsules at the outset of the task (similar to partial exposure tasks), they may have felt that they ‘knew’ what the contents of the selected capsule were. These alternative explanations may go some way to explaining why some younger children appeared to have understood the outcome of the demonstration but were not then able to use the information.

In line with our predictions, appropriate information use increased with age demonstrating that older children were better able to strategically adopt an appropriate response. This age-related difference may reflect younger children’s reduced capacity to represent the task goals and interpret the outcome of the demonstration. However, the performance of even the oldest children was not particularly high, suggesting that applying understanding of others’ goals relative to one’s own is challenging even if others’ goals are understood and the content of the social information has been correctly inferred. Again, in contrast to our predictions, we found no effect of the Same and Different Goals manipulation, suggesting that the degree to which the demonstrator’s and participant’s goals aligned did not influence children’s ability to strategically use social information. Although children were more likely to correctly infer the outcome of successful demonstrations, this did not extend to appropriate use of the information. Appropriate social information use was greater when the appropriate response was to shift, but not when the appropriate response was to copy. However, this pattern of results was also related to children’s age. It became clear that 3- and 4-year-olds’ appropriate information use did not follow the same pattern as that of the older children. While 3- and 4-year-olds’ performance remained no different to chance for both response types, 5-, 6-, and 7-year-olds’ performance was no different to chance for copying responses but was significantly above chance when the appropriate response was to shift. Therefore, the age-related improvement in children’s use of appropriate responses was restricted to trials in which the appropriate response was to shift. The chance performance of 3- and 4-year-olds may have been in part due to high cognitive demands taxing their executive functions such as memory. Indeed, we found they recalled which location the demonstrator had selected a capsule from less often than the older children. Below we discuss possible explanations for this pattern of performance.

Appropriate use of social information may indicate that children have recognised the value of the information that has been provided by others. Such reasoning-based understanding may suggest that children are employing explicitly metacognitive SLSs. If this is the case it would suggest that children can account for others' goals relative to their own, correctly interpret the content of social information, and apply this understanding to strategically adopt an appropriate response. Though the results suggest that children can account for others' conflicting motivations, the age-related improvement being restricted to appropriate shift responses suggests that some older children are not using explicitly metacognitive SLSs. Rather, they may have been relying upon implicit biases. The adult sample offers a useful comparison for what to expect from a participant who is using explicitly metacognitive social learning as we are confident that adults are capable of such a strategy. Therefore, the distinct difference between the performance of the oldest children and the adults supports our interpretation that some children were not using explicitly metacognitive strategies. Alternatively, this pattern of results may be an artefact of the binary choice task. Other tasks that have used a logically similar binary choice reward structure have also reported better performance in children when the appropriate response was to shift, albeit in slightly younger populations than the sample tested here (Atkinson et al., 2020). In a binary choice task, in which the demonstrator and participant share the same goal, a shift response following an unsuccessful demonstration performs an explore function which does not need to be traded off against an exploit option since the participant already knows that the alternative is undesirable. In contrast, appropriate copying following a successful demonstration involves a trade-off between explore and exploit motivations. When the demonstrator and participant have different goals, the opposite applies. Overall, this will inevitably favour shift responses (Gopnik, 2020). Evidence does suggest that children often choose to explore even in situations in which they might benefit from exploitation (Blanco & Sloutsky, 2020). Thus, it is possible that even in cases in which children recognise the value of the social information and can identify the appropriate response to achieve their goal, the motivation to explore may override the motivation to achieve the goal, making it difficult to ascertain the reason for particular decisions.

Previous research into SLSs has focused primarily on social information use with regards to copying. However, using tasks in which the demonstrated solution is always successful, or the desired response is to copy (e.g., Carr et al., 2015; van Leeuwen et al., 2018), makes it difficult to distinguish between preferences for exploration and a failure to copy. This means that researchers have sometimes concluded that children have failed to use the available social information to inform their responses when an alternative solution has been employed. Recent evidence suggests that in these cases children could be driven by a preference for exploring alternative options rather than simply copying or exploiting the options previously presented to them (e.g., Atkinson et al., 2020). Children who do not copy may in fact be fully cognisant of the social (or asocial, Blanco & Sloutsky, 2020) information available (e.g., potentially perfectly capable of producing the demonstrated response) but applying this with a motivation to explore. Indeed, in the current study the inclusion of both successful and unsuccessful demonstrations highlighted a distinct difference in children's performance depending on whether the appropriate response was to copy or avoid copying. Older children performed better at our task, but had we only included successful demonstrations involving aligned goals, this would not have been apparent.

We suggest that this motivation to explore should also be distinguished from favouring information acquired through personal experience over that acquired from a social source, which has been a common interpretation of social learning studies which have found low rates of copying. In the context of a social learning task, explore responses could reflect a preference for using personally-acquired over socially-acquired information. However, in the absence of comparison with an individual learning condition, it is impossible to distinguish between a preference for personal (over social) information from a preference for exploring new (over known) options. Indeed, recent research shows that children's preference for exploration is not peculiar to contexts of social information use; it was applied just as strongly following personally-acquired and social information (Atkinson et al., 2020; Renner et al., 2021). Similarly, Blanco and Sloutsky (2020) found a preference for exploration in an individual learning task even when children might benefit from exploitation. Such findings suggest that explore responses could indicate that children are motivated by a preference for new information. The results of our study are consistent with this interpretation.

Our discovery of age-related differences between children's goal understanding, demonstration outcome understanding, and task performance suggests 'jumps' or 'step-changes' that may be due to the development of abilities required for the use of different types of selective social learning. Firstly, there appears to be a jump from correctly imputing the goal to accurately inferring the outcome of the demonstration. This suggests that children (in this sample, around age 4 and 5) may be able to identify their own and others' goals, but not yet use that goal understanding in addition to the demonstrator's reaction to infer the outcome of the demonstration. There also appears to be a second jump from children's understanding of the demonstration outcome to being able to apply this understanding to inform their choice of response. This likely involves representing others' goals relative to one's own while also considering the outcome of the demonstrator's selection and recognising that the information provided by the demonstrator is valuable for one's own selection. The poor rate of explicit verbal reasoning indicates that children found articulating justifications for their choices during the task particularly challenging. Again, this is consistent with our suggestion that many children may have relied on implicit SLSs rather than being driven by reasoned understanding of the value of the social information.

Overall, these findings suggest that children are capable of understanding and representing others' goals from around 4 years whether they are aligned or not with their own. We also demonstrated that children younger than 6 years find it difficult to infer, from others' reactions, the content of social information when the outcome, or reward value, is opaque. Older children's more consistent performance was likely underpinned by reasoning-based understanding of others' access to information. Four- and five-year-olds' difficulty did not appear to be related to understanding others' goals. Rather, better performance following successful demonstrations indicated that some children were perhaps relying on less cognitively demanding biases based on more salient cues. In addition to this, it appears that using information about others' goals, whether they are aligned to one's own or not, poses further challenges. So, even if children have the ability to interpret the content of social information, they may still struggle to recognise its value to themselves. Such reasoning is likely dependent on the capacity for explicitly metacognitive SLSs which are proposed to be experience-dependent and to develop relatively late in childhood (Blakey et al., 2020, Chapter 3 of this thesis). This late development may go some way to explaining the below-ceiling performance of even the older children in using social information appropriately. In the absence of the capacity for explicitly metacognitive SLSs younger children's choices may have instead been determined by implicit biases and motivations towards exploration. This developmental trajectory highlights the significant cognitive challenge associated with using social information strategically and suggests that such abilities may be rare, or even absent, in animals. As cumulative culture may depend on strategic social information use in relation to specific goals, this late development is consistent with accounts proposing that such abilities may be linked to the distinctiveness of human cumulative culture.

Chapter 3:

Development of Strategic Social Information Seeking: Implications for Cumulative Culture

---

Chapter two demonstrated that children can account for others' conflicting motivations from around 4 years and they can reliably interpret the outcome of others' behaviour from around 6 years. While younger children's performance suggested that they might have been relying on implicit SLSs to correctly interpret the content of social information, older children's better performance appeared to reflect an understanding of the information's value indicating the capacity to employ explicitly metacognitive SLSs. However, appropriate use of information based on inferences about others' goals appeared to be challenging. The age-related improvements in appropriate information use applied only when the correct response was to shift, suggesting that even the older children may not have been using explicitly metacognitive SLSs.

While children's use of social information can indicate to some degree their understanding of it, looking at how they seek out or select information to attend to offers more potential in terms of gaining an insight into their recognition of its value. Indeed, in reality, learners may very rarely passively acquire social information that is valuable and relevant for achieving a specific goal. More likely, they have to seek out information from another individual who is in a position to provide the required information. Therefore, in a bid to better understand the developmental trajectory of explicitly metacognitive SLSs, the study reported in chapter three considers children's ability to seek out information based on an understanding of what information they require and who can provide that information.

The following chapter has been published as a preprint on PsyArXiv and is under revision for submission for publication, the references are given below. The chapter is presented in its submitted form. Data and analysis code are available in the OSF repository: [https://osf.io/ctse7/].

Blakey, K.H., Rafetseder, E., Atkinson, M., Renner, E., Cowan-Forsythe, F., Sati, S.J., and Caldwell, C.A. (Under Revision). Development of strategic social information seeking: Implications for cumulative culture.

Blakey, K.H., Rafetseder, E., Atkinson, M., Renner, E., Cowan-Forsythe, F., Sati, S.J., and Caldwell, C.A. (2020). Development of strategic social information seeking: Implications for cumulative culture. *PsyArXiv.* https://doi.org/10.31234/osf.io/4envx

Abstract

Human learners are rarely the passive recipients of valuable social information. Rather, learners usually have to actively seek out information from a variety of potential others to determine who is in a position to provide useful information. Yet, the majority of developmental social learning paradigms do not address participants' ability to seek out information for themselves. To investigate age-related changes in children's ability to seek out appropriate social information, 3- to 8-year-olds (*N* = 218; recruited in a school and at a public visitor attraction) were presented with a task requiring them to identify which of four possible demonstrators could provide critical information for unlocking a box. Appropriate information seeking improved significantly with age. The particularly high performance of 7- and 8-year-olds was consistent with the expectation that older children's increased metacognitive understanding would allow them to identify appropriate information sources. Appropriate social information seeking may have been overlooked as a significant cognitive challenge involved in fully benefiting from others' knowledge, potentially influencing understanding of the phylogenetic distribution of cumulative culture.

Introduction

Seeking out relevant information from appropriate social sources is ubiquitous in human adults. Human adults may therefore demonstrate key differences in the way they seek, attend to, and use social information compared to children and non-human animals (henceforth animals). This propensity for identifying and gathering relevant social information has been proposed as one of a suite of cognitive mechanisms that may be required for distinctively human cumulative culture (Baldwin & Moses, 1996; Heyes, 2018a; Tomasello et al., 1993). By 'distinctively human' cumulative culture we are referring to the accumulation of beneficial modifications to cultural traits over successive generations of learners, which results in increased functionality or efficiency (Caldwell & Millen, 2008b; Mesoudi & Thornton, 2018; Tennie et al., 2009). Its rarity in animals and apparent importance in accounting for human evolutionary success has prompted interest regarding the emergence and development of cognitive mechanisms in human children that support cumulative culture (Caldwell et al., 2020; Dean et al., 2014; Mesoudi & Thornton, 2018). It may be the lack, limited scope, or inflexibility of mechanisms such as information seeking that impedes development of human-like cumulative culture in animals despite them showing evidence of culture (Allen, 2019; Aplin et al., 2015; Hobaiter et al., 2014) and social learning abilities (Hoppitt & Laland, 2013; Whiten, 2017b). The current study aimed to address the gap in our understanding regarding age-related changes in children's ability to seek out and use appropriate social information. Documenting the developmental trajectory of this capacity could provide insights into the cognitive demands involved and whether it is likely to be observed in animals.

In a bid to understand the discontinuity between humans' and animals' capacity for cumulative culture a number of sociocognitive mechanisms have been considered (Boyd & Richerson, 1996; Dean et al., 2014; Lewis & Laland, 2012; Tennie et al., 2009). Some of these proposals focus on when and how social information is used. Social learning is generally considered to be adaptive when it is present within a population. However, models have shown that under most conditions it is the ability to flexibly switch from social to individual learning, when social learning proves unsatisfactory, that increases adaptability in a population (Enquist et al., 2007; Henrich & McElreath, 2003; Rendell et al., 2011; Rogers, 1988). That is, for social learning to be most beneficial and more efficient than individual learning it must be used selectively. Selective or flexible rules that influence individuals' use of social information are referred to as social learning strategies (SLSs). These refer to heuristic biases or rules that dictate when, what, and from whom social information should be acquired (Kendal et al., 2018; Laland, 2004), helping to filter out less useful aspects of available social information. The selective nature of SLSs such as 'copy older individuals' or 'copy the majority' makes them generally more effective than learning indiscriminately from others or individual learning. However, extensive evidence of SLSs in both young children and animals (Horner et al., 2010; Jiménez & Mesoudi, 2019; Laland, 2004; Price et al., 2017; Wood et al., 2012) suggests that selective social learning cannot account for the marked differences we see between humans and animals with regard to the capacity for cumulative culture.

One proposal that attempts to explain the distinctiveness of human cumulative culture outlines a dual-process view of social learning. This view suggests that there are distinct categories of SLSs each based on specific types of decision rules (Heyes, 2016c). According to this account, distinctively human cumulative culture could be attributed to the use of explicitly metacognitive SLSs (Dunstone & Caldwell, 2018; Heyes, 2016c). Explicitly metacognitive SLSs are defined as consciously represented and reportable rules. This means that learners are aware of the reasoning-based learning strategies that they are employing, and may also explicitly recognise states of knowledge, ignorance, and uncertainty in both themselves and others. Such strategies involve the use of theory of mind and metacognitive processes that enable human adults to flexibly identify, select, or disregard social information across varied contexts (De Oliveira et al., 2019; Dunstone & Caldwell, 2018; Heyes, 2016c, 2018a). As such, explicitly metacognitive SLSs afford learners the capacity to reason about what information is required and recognise the potential value to themselves of information that can be provided by others. These strategies are proposed to be experience-dependent and learned through social interaction, therefore we would expect them to emerge relatively late in development (Heyes, 2016c).

The second part of Heyes's (2016c) dual-process account of social learning proposes that the majority of the behaviours that conform to SLSs in both humans and animals are based on decision rules that depend on general-purpose associative learning processes or biologically selected biases. These rules direct learning towards objects, agents, and events that are most likely to provide useful information. However, unlike explicitly metacognitive SLSs, these *implicit SLSs* (as we will refer to them here) are not driven by a causal understanding of the potential value of social information. Although we refer to these strategies as implicit, we do not claim that learners employing them are necessarily devoid of insight regarding personal preferences that guide their social learning. It is likely that learners sometimes explicitly represent strategies related to salient yet superficial cues without appreciating *why* their strategy is successful. However, we suggest that such strategies, while they may be explicitly represented, are not explicitly metacognitive due to the absence of a causal understanding of informants' potential to provide valuable information. Social learning biases in young children and animals are likely driven by such implicit, and relatively crude, heuristic decision rules (Heyes, 2017). Indeed, cases in which very young children (Jaswal & Neely, 2006; Koenig & Harris, 2005; McGuigan, 2013; Rakoczy et al., 2010) and animals (Horner et al., 2010; Kendal et al., 2015; Ottoni et al., 2005; Price et al., 2017) select knowledgeable others are likely the result of implicit SLSs (e.g., associative learning or biologically selected biases), rather than the reasoning-based strategies employed by adults. For example, model-based biases for older (Rakoczy et al., 2010; Wood et al., 2012; Zmyj & Seehagen, 2013), higher-status (Jiménez & Mesoudi, 2019; McGuigan, 2013), or even more reliable models (Koenig & Harris, 2005) may be the result of repeated exposure to the successes of models with these characteristics, resulting in rule-like strategies.

While there is evidence of implicit SLSs, based on heuristic biases, in both young children and animals, in our view, there is as yet no solid evidence of explicitly metacognitive SLSs in either population. By contrast, adult humans are (with good reason) assumed to be able to use social information in an explicitly metacognitive manner (although this should not be taken to mean that they do not also use implicit SLSs). For example, it is routine for human adults to actively *seek out* models of social behaviour, using their understanding of others' intentions relative to their own to strategically select and use the most relevant social information (Caldwell et al., 2018; Vélez & Gweon, 2019). A critical question, then, is how children's use of social information develops with age, particularly in relation to their explicit understanding of the value of social information, and how they might be able to benefit from it. Since we assume that this understanding is not present from infancy, but is certainly in place by adulthood, it follows that this transition must occur over the course of childhood. We may therefore be able to identify key stages during development when children begin to change how they respond to social information (i.e., developing the ability to reason about its value). That is, it might be possible to identify when children develop the capacity to employ explicitly metacognitive SLSs that are driven by the value of the information, rather than relying on implicit SLSs (e.g., pre-existing biases from personal experience or phylogenetic history). This would provide insights into the cognitive capacities upon which such abilities depend, and therefore might also shed light on the reasons for the apparent absence of these abilities in animals.

Research into social learning has largely been restricted to investigating the circumstances under which social information is used, the efficacy of that use, and its role in cultural transmission (Burdett et al., 2016; Evans et al., 2018; Hoehl et al., 2019; Over & Carpenter, 2012; Rawlings et al., 2017; Subiaul et al., 2016; Wood et al., 2013a). Indeed, to date, developmental research into SLSs has focused on examining children’s responses to task-relevant social information (usually an effective solution) which is provided in advance of an opportunity to solve the same task (Dean et al., 2012; Horner & Whiten, 2007; Lucas et al., 2017; van Leeuwen et al., 2018; Wood et al., 2013a). However, while cultural transmission necessitates using information acquired from a social source, human learners are rarely passive recipients of valuable social information. Here, we propose that *actively* seeking out valuable information when faced with a particular problem to solve is more analogous to real world social learning scenarios, compared with being *passively* provided with relevant information.

Therefore, in the current study we were particularly interested in determining when children develop the ability to seek out social information based on an understanding of the value of that information. This led us to consider selective information seeking paradigms that already exist in the literature. Specifically, we looked at the selective trust paradigm. This paradigm is commonly used to examine children’s preferences for information provided by models with conflicting social and/or epistemic characteristics, and is built on the premise that children learn from others’ testimony (see Harris et al., 2018 for a full review).

These studies typically employ a ‘conflicting sources paradigm’ in which children first observe two informants who differ on social (e.g., gender or accent) and/or epistemic (e.g., accuracy or reliability) characteristics (Koenig & Harris, 2005). Following a familiarisation phase, children are faced with an unfamiliar scenario, for example, naming an unfamiliar object. In some tasks, children are required to select one of two potential informants to ‘seek’ information from (i.e., ‘ask’ questions) and/or required to make their selection following the informants’ claims about the name of the object (i.e., ‘endorse’ questions). The model selected by the child is considered to be the model whose claim they trust. With regards to the influence of social characteristics on selective trust, evidence suggests that in the absence of epistemic differences children ask and endorse informants who have positive social characteristics (Tong et al., 2020). These selections are influenced by both the models’ and the learners’ own characteristics (e.g., age or gender, Wood et al., 2013b). In particular, social characteristics that signal a model’s similarity to the learner (i.e., ingroup membership) are consistently favoured by young children (e.g., preference for informant with a native accent, Corriveau et al., 2013; preference for informant of the same gender, Terrier et al., 2016). Model-based biases such as these are unlikely to require explicit cognitive reasoning. Similar to the selective preferences in children’s proclivity to copy and consistent with the SLS ‘copy successful individuals’, they are likely the result of implicit biases that promote learning from sources that are most likely to provide useful information across the broadest range of contexts. The selective trust literature reports that children as young as 3 years old are sensitive to informants’ social and epistemic characteristics (Corriveau et al., 2013; Jaswal & Neely, 2006; Koenig & Harris, 2005; Sobel & Kushnir, 2013). Recent meta-analyses examined the relation between children’s age and their selective trust decisions (Tong et al., 2020). Results indicated that children asked and endorsed more knowledgeable (accurate/reliable) informants when they differed on only epistemic characteristics. Specifically, 4-year-olds were more likely than 3-year-olds to endorse knowledgeable informants. When informants differed on epistemic and social characteristics simultaneously, 4- to 6-year-olds were more likely to endorse informants who were knowledgeable but had a negative social characteristic while 3-year-olds appeared to weigh both characteristics equally. Thus, from 4-years children appear to place greater value on epistemic characteristics.

Whilst these preferences may be linked to a developing explicit awareness of the potential value of social information, the design of these studies precludes this conclusion. The paradigm depends on participants being exposed to information about the accuracy or reliability of the two informants in the familiarisation phase, in order to establish the respective epistemic characteristics (e.g., knowledgeable/ignorant) of the conflicting sources. Thus, the literature on selective trust, like much of the literature on SLSs, depends on children making choices between models on the basis of characteristics for which they are likely to have a prior history of associations or a pre-existing bias (whether established as part of the experimental procedure, e.g., reputation for accuracy, or from the child's own life experience, e.g., age, VanderBorght & Jaswal, 2009). The results can therefore be likewise attributed to implicit biases. However, the apparent transition to favouring epistemic characteristics (such as prior accuracy) over social ones (such as familiarity, Corriveau & Harris, 2009) is perhaps suggestive of a developing insight into the value of others as sources of social information.

While it is useful to know how children use social information and who children prefer to learn from, we argue that this is not sufficient to determine the cognitive mechanisms that children are employing during social learning. Children's ability and proclivity to seek out *relevant* social information has not yet been adequately addressed. That is, nothing in the SLSs or selective trust literature has examined children's ability to select valuable social information on the basis of its relevance for solving a specific problem.

Explicitly metacognitive SLSs are proposed to depend on relatively late developing cognitive capacities (Heyes, 2016c). Therefore, to help explain why cumulative culture appears to be restricted to humans we can look at whether age-related changes in children's ability to seek out relevant social information coincide with advances in cognitive development. If we find particular ages at which children make significant advances in their appropriate social information seeking and these occur at a similar age to the development of particular cognitive capacities, then these capacities might be necessary prerequisites. Similarly to Heyes (2016c), Baldwin and Moses (1996) proposed that motivations for initiating an appropriate search for social information likely rest on advanced metacognitive capacities. They emphasised that to seek social information effectively, the seeker should have awareness of what information is required and from whom it can be obtained. While behavioural tests of implicit metacognitive ability suggest that infants (Goupil et al., 2016), young children (Bernard et al., 2015), and animals (Call & Carpenter, 2001; Neldner et al., 2015) react to the state of ignorance, they do so without necessarily recognising a metacognitive awareness of that state. However, we know that children's cognitive capacities continue to progress well beyond those of animals. In particular, as their cognition advances, children develop abilities such as evaluating their own knowledge state (Kloo & Sodian, 2017; Rohwer et al., 2012), understanding others' mental states (Krachun et al., 2009; Wimmer & Perner, 1983), and recognition that perceptual access to information facilitates knowledge formation (O'Neill et al., 1992; Pillow & Weed, 1997; Ruffman & Olson, 1989). Such abilities are thought to be cognitively demanding; thus, their requisite nature may preclude younger children's (and by logical extension animals') ability to seek out appropriate social information. Therefore, if we can identify when children develop the capacity to use explicitly metacognitive over implicit SLSs we may be able to use this to predict the likelihood of the capacity being available to animals. That is, if we find that this ability develops late in childhood then this would be consistent with the hypothesis that it is dependent on cognitive capacities that are not available to animals, and could help to explain the distinctiveness of human cumulative culture.

The aim of the current study was to investigate when children develop the ability to seek out social information using explicitly metacognitive SLSs. To explore the development of children's appropriate information seeking we designed a task in which 3- to 8-year-old children were faced with a problem for which they could not use prior experience or knowledge. Rather, to solve the problem children had to reason about the information needed and who had the potential to provide that information. That is, appropriate information seeking should be based on a reasoned understanding of the value, to themselves, of the social information. Heyes (2016c) argued that this kind of cognitive reasoning may not be available in young children or animals. In the absence of the ability to reason about the potential value of social information provided by others, we expected children to rely on less cognitively demanding implicit SLSs. In the current study we specifically looked at whether children's information seeking might instead be influenced by heuristic model-based biases for superficial demonstrator characteristics, such as age or gender. Thus, the task was designed such that the use of model-based biases in the absence of cognitive reasoning would be inappropriate, leading to imperfect information seeking.

Overall, we expected to find an age-related transition from the use of heuristic model-based biases to reasoning-based choices driven by the value of the information. Specifically, we predicted that appropriate information seeking would improve with age due to our anticipation that advances in children's metacognitive understanding would enable them to identify appropriate sources of information. We expected that younger children would struggle to employ explicitly metacognitive SLSs and instead rely on less cognitively demanding heuristic biases (implicit SLSs) related to superficial demonstrator characteristics such as age or gender. Any age-related changes in children's appropriate information seeking could indicate use of different SLSs when approaching the task. Thus, we examined the developmental trajectory of appropriate social information seeking to provide insight into the emergence of cognitive reasoning as a mechanism required for distinctively human cumulative culture. Finding evidence of reasoning-based choices in older, but not younger, children would support Heyes's (2016c) proposal that explicitly metacognitive SLSs are experience-dependent, developing relatively late in children.

Appropriate use of social information (following appropriate information seeking) could further help to distinguish between learners with some level of metacognitive understanding, and those reliant on implicit rules for what, when, and whom to copy. That is, understanding the relevance of the acquired information would be expected to also result in more appropriate *use* of that information. By investigating social information use in conjunction with appropriate information seeking we can really begin to expose the cognitive mechanisms that underlie these processes. How children use particular types of information provided by others can tell us much about the nature of their social learning processes – most importantly whether they might be using reasoning-based strategies or relying upon adaptive heuristic biases. In the current study, copying was not always the correct response, and success was sometimes dependent on making a different choice to the one made by the demosntrator. As with appropriate information seeking, we expected successful use of the social information to improve with age.

Method

Participants

The final participant sample comprised 218 children aged three to eight years (112 females; Mean age = 71.1 months, *SD* = 20.4 months, range = 37 to 106 months). The sample size was balanced across age groups: 3 years (*n* = 36), 4 years (*n* = 40), 5 years (*n* = 37), 6 years (*n* = 36), 7 years (*n* = 35), 8 years (*n* = 34). The sample includes children recruited from a school in Scotland (*n* = 118) and visitors recruited at Edinburgh Zoo (*n* = 100). The nationality of the sample, as identified by parents or legal guardians, was predominantly British. An additional eight children were tested but later excluded from analyses due to researcher or technical errors (*n* = 4), missing data (*n* = 2), and task interference (*n* = 2). This study was approved by the University of Stirling General University Ethics Panel (GUEP555), and informed written consent was provided by each participant's parent or legal guardian.

Materials

The task apparatus is presented in Figure 1. Four identical wooden boxes ($16cm^3$) were locked using child friendly plastic padlocks (Alphabet Learning Locks, Lakeshore, Carson, CA, USA). Each target padlock was a different colour (red, blue, orange, purple) with a different cartoon image (elephant, queen, pig, sun) on the front. Additional padlocks of different colours and with different cartoon images were used for non-target social demonstration videos. One locked box was presented in each trial along with two plastic keys. Only one key could unlock the padlock. A coloured shape (star, rectangle, circle, oval) was fixed to each key: within a trial the same shape was fixed to both keys each in a different colour. So that the colours of the keys were salient in the demonstration videos each key was placed on a rectangular laminated card of the same colour as the shape on each key. Each box contained a coloured rubber duck. Retrieved ducks were placed on a laminated image of a pond. A Lenovo Yoga 520 touchscreen laptop running PsychoPy v1.84.2 (Peirce et al., 2019) in tablet mode was used to present visual stimuli and videos.

Figure 1

Example of the Experimental Set-Up for One of the Four Trials

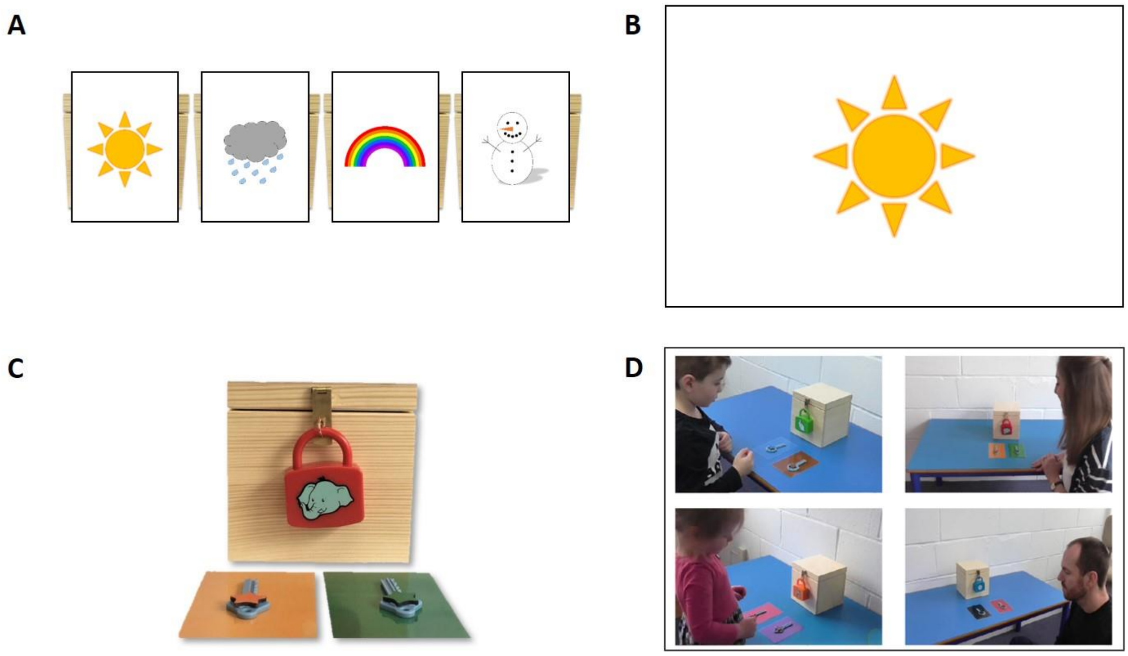

Note. A: trial indicator cards placed in front of boxes; B: trial indicator image displayed on screen; C: box locked with a distinctive padlock, and two possible keys (one correct); D: still images presented on the screen representing available demonstration videos, in this example the target video is on the top right.

Social Demonstration Videos

A total of 20 demonstration videos were created, including two possible target videos (successful and unsuccessful) and three non-target videos for each of the four trials. Each video depicted a demonstrator faced with a similar set-up to the participant, selecting one of two keys to try to unlock a padlock to open a box. Within each trial the age and gender of the four potential demonstrators was different (adult male, adult female, child male, child female; child demonstrators were all between 3 and 4 years) and each demonstrator had a different combination of padlock and keys from one another. In each trial a single target video showed a combination of padlock and keys which matched the participant’s combination of padlock and keys, and three non-target videos in which none of the apparatus matched that of the participant. Thus, the target video would provide useful information regarding which key would unlock the padlock, and the non-target videos would provide irrelevant information. The age and gender of the target demonstrator was different in each trial and the order of the trials was randomised between participants. The arrangement of the target and non-target demonstrators on the screen was also randomly assigned between trials and participants. A successful demonstration showed the selection of the key that unlocked the padlock and the box being opened, while an unsuccessful demonstration showed the selection of the key that did not unlock the padlock and the box remaining closed. Target videos were randomly allocated as successful or unsuccessful in each trial so that a participant would see a maximum of two videos of each demonstration type if they selected all four target videos. The three non-target videos in each trial comprised both successful and unsuccessful demonstrations.

Procedure

Participants took part individually in a single session that lasted approximately 10 minutes. Two experimenters were present, experimenter one (E1) provided instructions and presented materials to participants (see Appendix B for a verbal script), while a second experimenter controlled the laptop and live coded participants' responses. In some sessions, a familiar adult was also present but was asked not to interact with the participant during the task. Participants received a sticker for taking part.

At the outset each of the four locked boxes were placed behind one of four trial indicator cards (sun, rain, rainbow, snow) that corresponded to the trial in which that box was to be attempted (see Figure 1A). The participant touched the screen to begin the task, generating one of the trial indicator images (see Figure 1B) and allowing E1 to retrieve the corresponding box and keys for that trial from behind the trial indicator card. The box was brought to the front of the testing table and the two keys were placed on the correspondingly coloured cards side by side (see Figure 1C) between the box and the participant (keys were presented on the same sides as in the associated target demonstration video). E1 explained that one of the keys would unlock the padlock so the box could be opened and one of the keys would not unlock the padlock so the box would stay closed. Participants were told that they would watch a video before trying to open their own box. E1 explained that participants would *“see four pictures of other people who have boxes they want to unlock”* and that they could *“choose one of the pictures to see a video of that person trying to open their box”*. Participants were instructed to *“choose a video that will help you choose which key to try and open your box with”*. The laptop screen was turned away during the instructions to reduce distraction. When it was turned to face the participant, four still images from the beginning of each demonstration video were presented on the screen (one target and three non-target; see Figure 1D). Participants were asked *“Which video are you going to choose?”* to prompt them to select one of the videos to watch. Once an image had been selected, the corresponding demonstration video played immediately. When the video ended the screen turned white and was turned away; participants were told, *“Now it’s your turn to try and unlock your box”* and asked, *“Which key are you going to choose?”* If participants chose the target key and successfully opened the box, they retrieved the rubber duck. If participants chose the non-target key the box remained locked. The same procedure was repeated for each of the four trials. Following completion of all trials E1 asked *“How were you deciding which videos to watch?”* and *“How were you deciding which keys to use?”* to probe whether children were able to explicitly and reasonably justify their choices of videos and keys.

Statistical Analysis

The analyses were performed in R (R Core Team, 2020), with generalised linear mixed effects analyses (GLMMs) performed using the *lme4* package (Bates et al., 2014) with logit regression. P-values < .05 were accepted as statistically significant. The binary dependent variables in the analyses were: appropriate information seeking and appropriate information use. Where specified as fixed effects the following variables were sum coded: participant’s gender (female as −1, male as 1), participant-target demonstrator gender congruence (incongruent as −1, congruent as 1), target demonstrator age (child as −1, adult as 1) and demonstration outcome (unsuccessful as −1, successful as 1). Age was centred and scaled to measure thousands of days. The random effects structure for each model aimed to include by-participant random slopes for all fixed effects and keep random effects structures ‘maximal’ where possible (following Barr et al., 2013). Where the ‘maximal’ model resulted in non-convergent or singular fit models, random slopes were removed, followed by random intercepts where necessary, until a convergent, non-singular model was obtained. Post hoc analyses were carried out using estimated marginal means using the *emmeans* package (Lenth et al., 2019). Post hoc results are given on the log odds ratio scale.

Results

The two key aims of this study were to examine whether children were able to seek out and select the target demonstration video and, if successful, whether they were able to use the social information in the demonstration to select the target key to unlock the padlock. We were also keen to explore the errors that children were making with regards to information seeking, specifically looking for evidence of any model-based biases for demonstrator characteristics. Finally, we examined children's responses to post-test questions probing their explicit reasoning.

Information Seeking

Appropriate information seeking was measured by the proportion of trials in which children selected the target video. This required children to reason about the information they needed to solve a problem (unlock their box) and which of the demonstrators could provide that information (which of the demonstrators were facing the same problem). If children were reasoning appropriately, they should have disregarded demonstrator characteristics such as age and gender. Overall children selected the target demonstration video in 65% of 872 trials.

A GLMM was built for information seeking success with fixed effects of age, gender, participant-target demonstrator gender congruence, target demonstrator's age, and interactions between these variables, a random intercept of trial number and a by-participant random slope for target demonstrator's age. A significant main effect of age (b = 2.72, SE = 0.41, z = 6.70, p < .001) indicated that children's appropriate information seeking improved with age (see Figure 2), with 7- and 8-year-old children successfully selecting the target video in 90% and 91% of trials, respectively. Binomial tests revealed that appropriate information seeking was significantly above what would be expected by chance (25%) in each age group (p < .001).

Figure 2

Proportion of Target and Non-Target Demonstration Video Selections by Age in Years

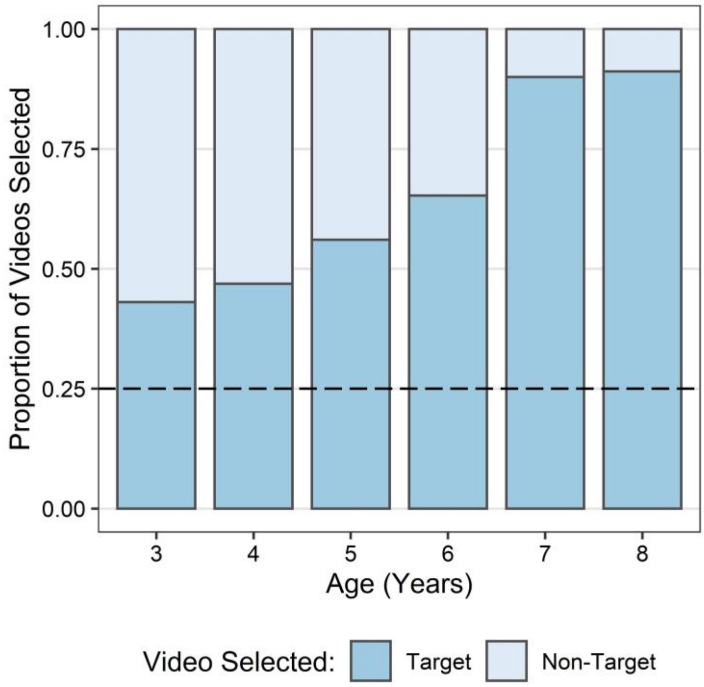

Note. Proportion of trials (*N* = 872) in which participants selected target and non-target demonstration videos. Target video selections equate to appropriate information seeking. All age groups selected the target video significantly above chance. Dashed line indicates chance.

A significant main effect of participant-target demonstrator gender congruence (b = 0.36, SE = 0.11, z = 3.22, p = .0013; see Figure 3) suggested that children selected the target video significantly more often when the gender of the target demonstrator was congruent (M = .69, SD = .46) rather than incongruent (M = .60, SD = .49) with their own. The GLMM did not reveal any evidence of an overall effect of gender (b = −0.15, SE = 0.19, z = −0.79, p = .429) or the target demonstrator's age (b = −0.09, SE = 0.18, z = −0.49, p = .622) on appropriate information seeking. However, there was a significant two-way interaction between gender congruence and participant gender (b = −0.23, SE = 0.11, z = -1.98, p = .048). To clarify this interaction, we performed a post hoc analysis using *emmeans.* This indicated that the gender congruence effect was present in female

(b = 1.17, SE = 0.33, z = 3.52, p = .0004) but not male participants (b = 0.27, SE = 0.30, z = 088, p = .378).

Figure 3

Proportion of Target and Non-Target Demonstration Video Selections by Participant-Selected Demonstrator Gender Congruence

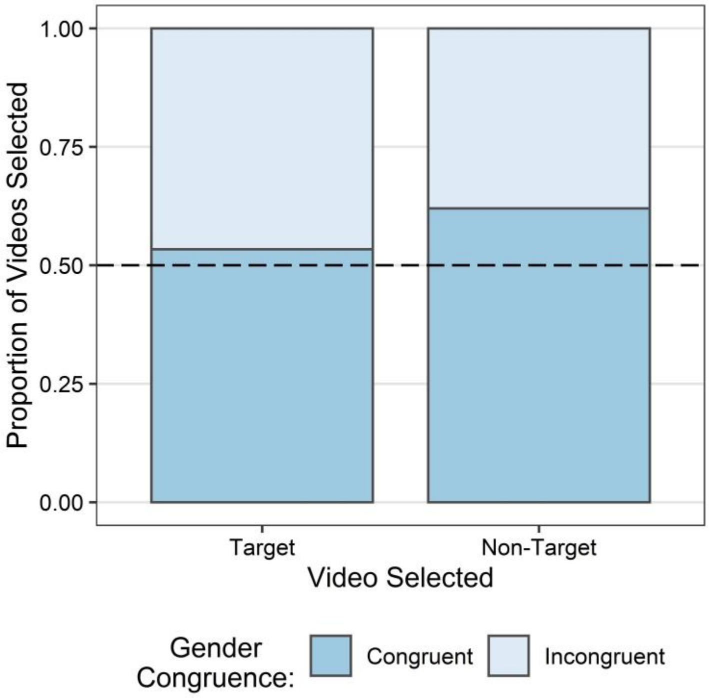

Note. Proportion of trials (N = 872) in which participants selected target and non-target demonstration videos by participant-selected demonstrator gender congruence. Target video selections equate to appropriate information seeking. Dashed line indicates chance.

A three-way interaction between gender congruence, target demonstrator age and participant gender was also found to be significant (b = 0.30, SE = 0.11, z = 2.63, p = .009). To clarify this interaction, we performed a post hoc analysis using *emmeans.* This indicated that the gender congruence effect found in female participants was restricted to cases in which the target demonstrator was a child (b = 1.60, SE = 0.68, z = 2.42, p = .016), while male participants showed a marginally non-significant gender congruence effect when the target demonstrator was an adult (b = 1.03, SE = 0.42, z = 2.47, p = .064). For an overview of the distribution of demonstrator selections split by participant age and gender see Figure 4.

Figure 4

Proportion of Selections of Each Demonstrator by Participant Age and Gender

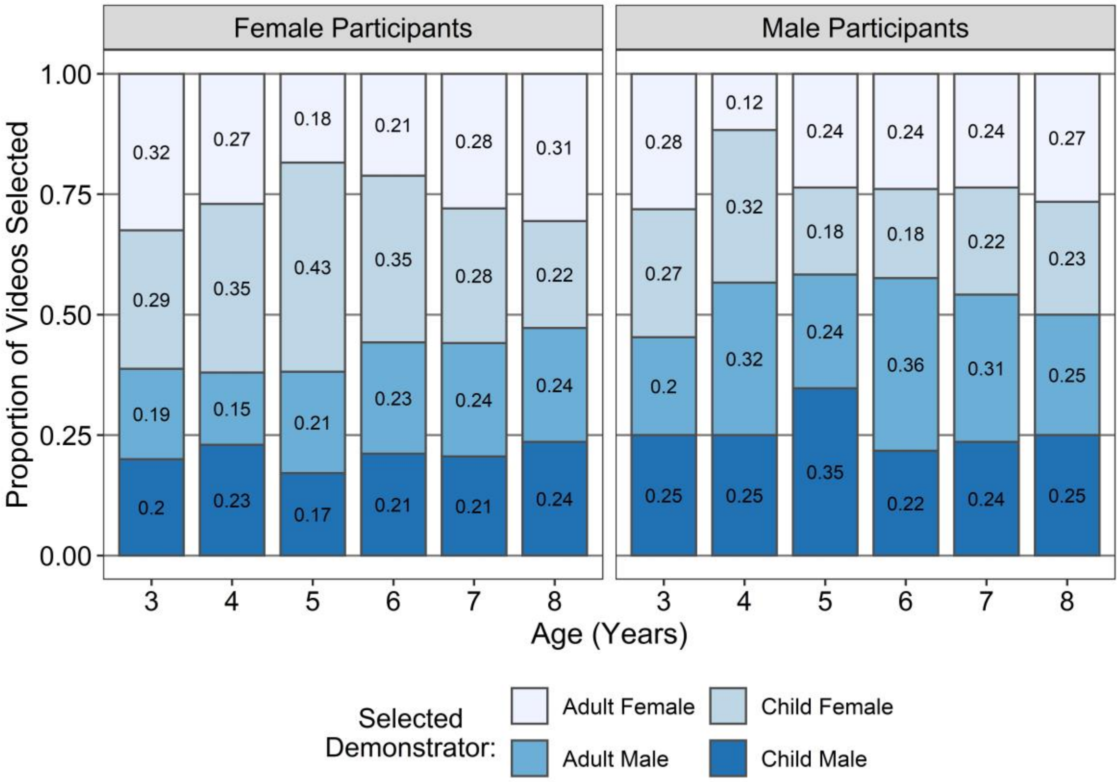

Note. Proportion of trials (N = 872) in which participants selected each demonstrator, split by participant age and gender. Optimal information seeking would have resulted in an equal distribution of selections of each demonstrator. Values presented on bars refer to proportion of selections in each section.

Information Seeking Errors

The proportion of trials in which children successfully selected the target demonstration video was significantly above chance in all age groups. However, in 35% of trials children selected a non-target video. Therefore, we assessed whether, in the absence of the ability to reason about the information needed, younger children were relying on heuristic model-based biases, pertaining to demonstrator characteristics. We focused on the errors of 3- to 6-year-old children due to 7- and 8-year-olds' high success rate (however inclusion of the 7- and 8-year-olds' errors does not alter the results).

To explore the information seeking errors we looked at which non-target demonstrators children were selecting relative to model-based characteristics, specifically demonstrator age (adult or child) or gender. A chi-square test of independence revealed no significant preference for either demonstrator age group (χ^2 (1) = 3.63, p = .057). However, a chi-square test of independence revealed a significant preference for selecting gender congruent over gender incongruent non-target demonstrators (χ^2 (1) = 14.53, p < .001; see Figure 3). A chi-square goodness-of-fit test was performed to compare the proportion of participants' selections of each demonstrator with the proportion that would be expected given equal distribution of selections across all four demonstrators. Goodness-of-fit results for female participants indicated a significant difference in the proportion of selections of each demonstrator (χ^2 (3) = 20.27, p < .001). Female participants showed a preference for female child demonstrators, selecting them significantly more often than chance (χ^2 (1) = 7.40, p = .007), while selecting the other demonstrators no differently to chance ($p \geq .14$) indicative of a specific model-based bias. Goodness-of-fit results for male participants indicated a distribution of demonstrator selections that was no different to chance (χ^2 (3) = 4.40, p = .22) suggesting that male children had no specific preference for any of the demonstrators.

Information Use

Appropriate information use was measured by the proportion of trials in which children, who had selected and watched the target demonstration video, selected the target key. The appropriate response was dependent on the success of the demonstration. Viewing a successful demonstration should have encouraged a copying response (using either the colour or the position of the key), while viewing an unsuccessful demonstration should have encouraged avoidance of that response, and selection of the alternative. Overall children chose the target key in 83% of 564 trials (287 successful and 277 unsuccessful demonstrations).

A GLMM was built for appropriate information use with fixed effects of age, demonstration outcome, and the interaction between these variables, and random intercepts of participant gender, target demonstrator (adult male, adult female, child male, child female), and participant ID. Similarly to the information seeking results, a significant main effect of age was revealed (*b* = 1.86, *SE* = 0.30, *z* = 6.26, *p* < .001) indicating that appropriate information use improved with age (see Figure 5). The model also identified a significant main effect of demonstration outcome (*b* = 0.69, *SE* = 0.18, *z* = 3.89, *p* < .001) indicating that children were more successful following successful demonstrations (*M* = .88, *SD* = .33), than unsuccessful demonstrations (*M* = .78, *SD* = .42). It is worth noting that appropriate information use was particularly high in 6-, 7-, and 8-year-old children following successful demonstrations, with children successfully copying the demonstrator's key choice in >95% of trials. Finally, there was a significant two-way interaction between age and demonstration outcome (*b* = 0.52, *SE* = 0.25, *z* = 2.10, *p* = .037). To clarify the direction of this interaction, we performed a post hoc analysis using *emmeans*. This indicated that the effect of demonstration success was slightly more pronounced in older children (upper quartile; *b* = 1.88, *SE* = 0.53, *z* = 3.52, *p* = .0004) than younger children (lower quartile; *b* = 0.85, *SE* = 0.29, *z* = 2.98, *p* = .003).

Figure 5

Proportion of Target and Non-Target Key Selections by Age in Years

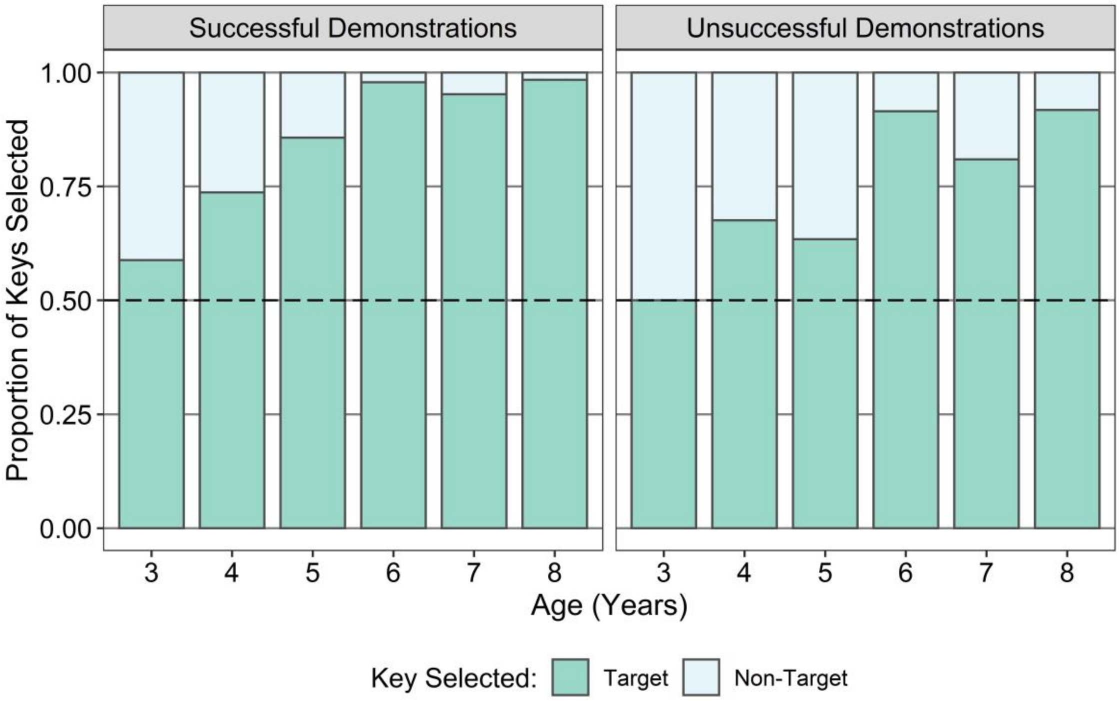

Note. Proportion of trials in which participants who watched the target demonstration video (*N* = 564) selected the target and non-target keys. Target key selections equate to appropriate information use. After successful demonstrations, the appropriate response was to copy the demonstrator's key selection; after unsuccessful demonstrations, the appropriate response was to avoid copying the demonstrator's selection. Dashed line indicates chance.

Explicit Verbal Reasoning

We examined children's responses to explicit reasoning questions regarding their choices in both the information seeking and information use aspects of the task. Responses to each reasoning question were categorised into four levels: 0) no response; 1) non-reasoned response – responses that did not relate to the question asked (e.g., *'the good ones'*/*'I just used my brain'*), or comprised single words or gestures; 2) reasoned but incorrect responses – explanations that showed evidence of explicit reasoning but the motivations were incorrect or the answer was insufficient to determine full correct reasoning (e.g., *'picked one of the grown ups'* /*'because some are wrong and some are* *right'*); and 3) reasoned correct response – the explanation provided clear reasoned evidence of explicit task understanding (e.g., *'because they had the same lock'/'in the video, if they didn't get it I chose the other one, if they did I chose the same one'*). When asked to justify how they were selecting which videos to watch, 38% of children provided correct reasoned responses, and when asked to justify how they were selecting which keys to use, 23% gave correctly reasoned responses.

We investigated whether children's verbal reasoning was predicted by age or task performance by conducting two ordinal regressions (using the ordinal package in R, Christensen, 2019). For each reasoning question (information seeking and information use) we submitted children's responses to an ordinal regression with fixed effects of age, performance (total number of target video and target key selections, respectively), and the interaction between these variables. Both models indicated significant main effects of both age (information seeking: b = 1.38, SE = 0.61, z = 2.24, p = .025; information use: b = 2.49, SE = 0.54, z = 4.61, p < .001; see Figure 6) and performance (information seeking: b = 0.59, SE = 0.12, z = 5.11, p < .001; information use: b = 0.39, SE = 0.11, z = 3.44, p < .001; see Figure 7). There was no evidence of any significant interactions between age and performance (information seeking: b = 0.24, SE = 0.21, z = 1.18, p = .237; information use: b = −0.07, SE = 0.19, z = -0.35, p = .729). These results suggest that in response to both reasoning questions older children and children who made more target selections provided better reasoned responses than younger children and children who made less target selections, respectively.

Figure 6

Proportion of Participants Providing Each Category of Explicit Verbal Reasoning Responses by Age in Years and Reasoning Question

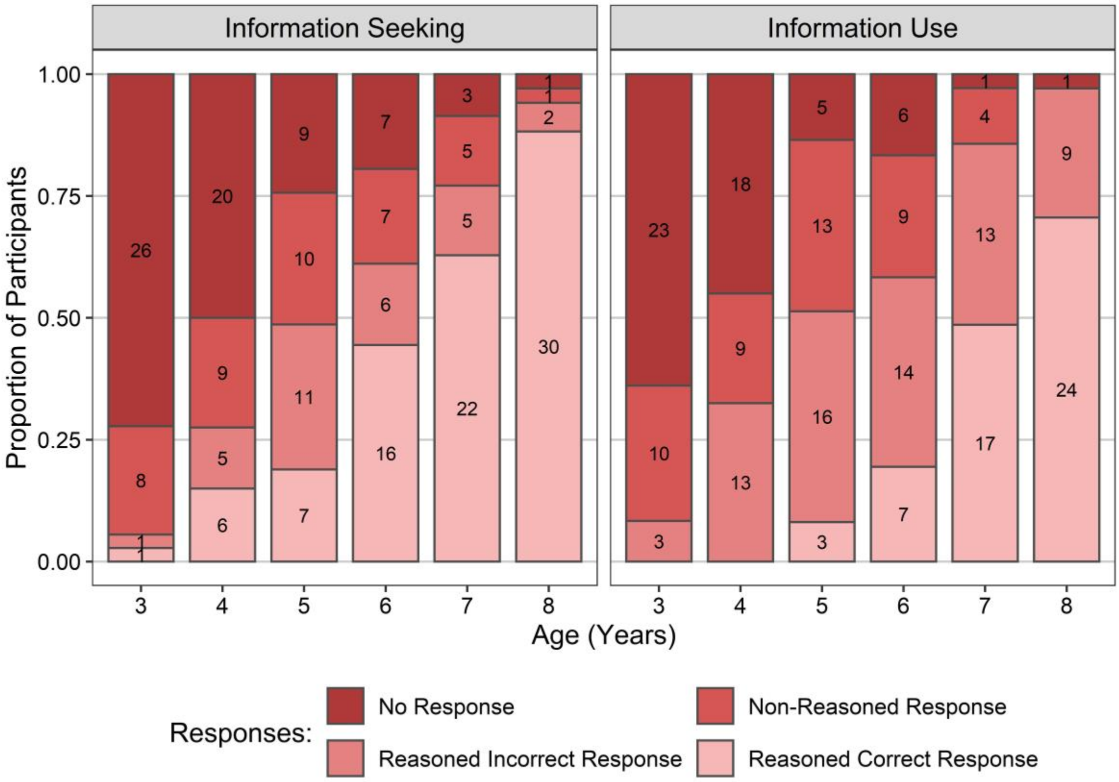

Note. Proportion of participants in each age group who provided each category of responses to the explicit reasoning questions about how they sought out and used the social information. All participants are included (*N* = 218) regardless of whether they sought out the correct information. Values presented on bars refer to the number of participants each proportion represents.

Figure 7

Proportion of Participants Providing Each Category of Explicit Verbal Reasoning Responses by Task Performance and Reasoning Question

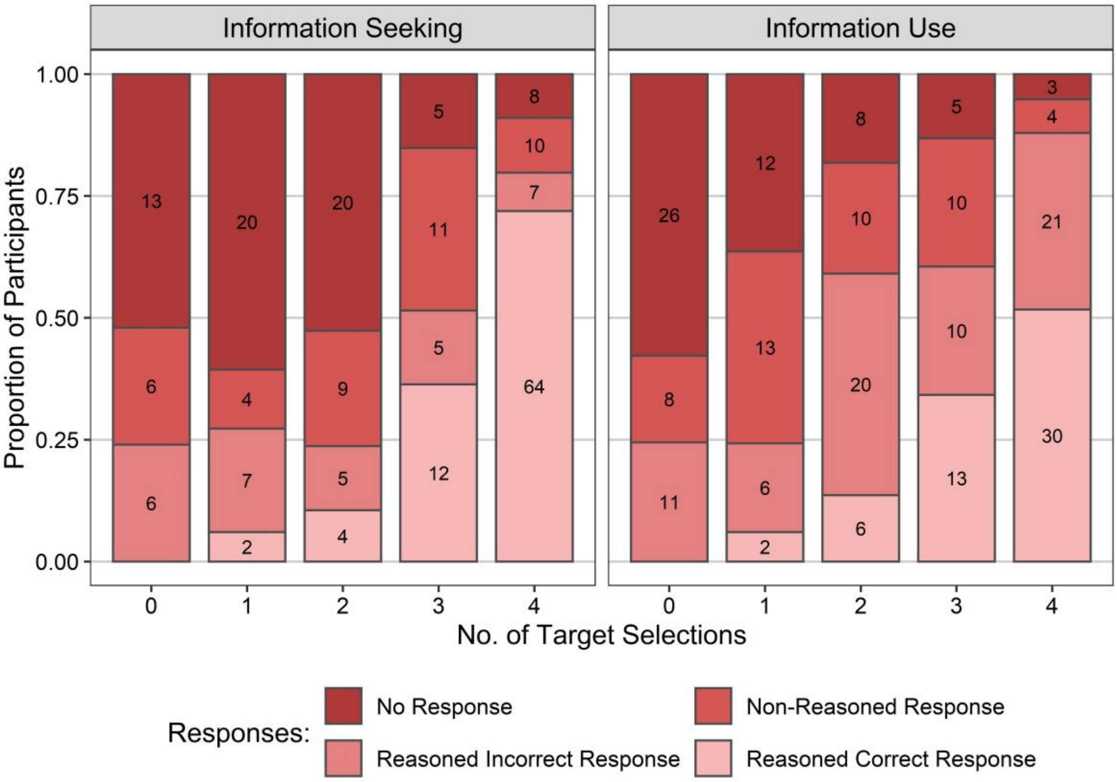

Note. Proportion of participants who provided each category of responses to the explicit reasoning questions about how they sought out and used the social information (*N* = 218). Information Seeking: responses by target video selections; Information Use: responses by target key selection. Values presented on bars refer to the number of participants each proportion represents.

Discussion

In this study we investigated the development of appropriate social information seeking in 3- to 8-year-olds. Specifically, we were interested in identifying when children develop the ability to seek out information based on a reasoned understanding of what information is required to solve a particular problem and who could provide it (explicitly metacognitive SLSs). We were also interested in exploring the alternative strategies children rely on prior to this development, therefore we examined the possible influence of model-based biases (implicit SLSs), which (if the task had been understood) should have been disregarded.

We found that the proportion of trials in which children selected the target demonstration video increased with age. This corresponded to our prediction that older children (7- to 8-years) would exhibit more appropriate social information seeking than younger children (3- to 6-years). Despite finding that all ages selected the target video significantly more often than would have been expected by chance, we saw relatively high information seeking error rates in the younger age groups. This suggests that older and younger children might have been approaching the task differently. The particularly high rate of appropriate information seeking we observed in 7- and 8-year-olds (close to ceiling) was consistent with our expectation that increased metacognitive understanding would allow them to recognise what information was required to solve the problem and to identify who, of the available demonstrators, could provide relevant information. Thus, we suggest that older children understood the potential value to themselves of the information that the target demonstrator could provide having employed an explicitly metacognitive SLS. By contrast, we postulate that younger children’s target video selections were less likely to have been driven by reasoned understanding of the information they needed and the potential value of the target demonstration video. Therefore, when selecting a video, younger children might have instead relied on less cognitively demanding implicit SLSs such as model-based biases, or a bias for selecting models whose box matched their own box (without appreciating the value of the information).

The primary aim of this study was to establish when children develop the ability to seek out information based on a reasoned understanding of its value. However, we were also interested in what strategies children used to seek out information in the absence of this ability. The design of this study made it possible to explore whether children might be relying on less cognitively demanding implicit SLSs such as heuristic model-based biases. If children understand the value of social information and its relevance to their goal, then it should not matter what age or gender the appropriate demonstrator is. That is, appropriate information seeking based on an understanding of the value of the information should override model-based biases. Therefore, in our analyses we assessed whether children's information seeking showed evidence of any biases towards either of the demonstrator characteristics we manipulated, namely age and gender. Results showed that children were more likely to select target demonstrators who were the same gender as themselves, therefore showing a gender congruence effect. Our exploratory analysis into the information seeking errors made by 3- to 6-year-olds revealed a similar effect of gender congruence. While there were slight differences in particular demonstrator preferences between male and female participants, we found no overall bias related to the demonstrators' age. Therefore, it seems that even when it was apparent – from salient visual cues – which demonstrator could provide relevant information, some children were still driven by biases for superficial model characteristics (in this case demonstrator gender) rather than the potential value of the information. These findings support the interpretation that in the absence of the capacity for explicit reasoning about what information is needed and who can provide it, younger children were influenced by model-based biases. These biases are in line with the reports of preferences for same-gender models in selective trust paradigms (Terrier et al., 2016). That the gender congruence effect extended across both target and non-target demonstrator selections suggests that some target demonstrator selections were probably made as a result of model-based biases rather than reasoning about the value of the information that the demonstrator could provide.

As outlined above, we propose that the older children's high rate of appropriate information seeking was driven by reasoned understanding of the value of the information, whereas the younger children who have likely not yet developed the capacity to use these reasoning-based SLSs instead relied on model-based biases (in this case a preference for gender congruent demonstrators). The results of the current study highlight the importance of considering cognitive mechanisms involved in seeking out social information, as they reveal evidence of the different strategies that older and younger children use when approaching the task. Thus, we believe this provides evidence of a relatively late age-related transition from implicit SLSs to explicitly metacognitive SLSs, consistent with accounts proposed by Dunstone and Caldwell (2018) and Heyes (2016c).

In line with our predictions, the proportion of trials in which children selected the target key (appropriate information use) increased with age. Rates of target key selections were found to be higher in older children following both successful and unsuccessful demonstrations, indicating that older children were more adept at flexibly using social information to inform their responses than younger children. We propose that older children knew what to do with the information they received in the demonstration because their appropriate information seeking had been based on an understanding of the information needed to solve their problem. Overall, these results provide further evidence for the interpretation that older children’s social information seeking was based upon reasoned understanding of the value of the information (explicitly metacognitive SLSs). By contrast, younger children appeared to make less effective use of the information they acquired despite selecting the target video. We posit that this offers further evidence to support our interpretation that younger children’s target video selections were less likely to have been driven by reasoned understanding of what information was required and the potential value of the target demonstration video. Rather, their selections might have been based on superficial but salient demonstrator characteristics driven by implicit SLSs (in this case, model-based biases for gender congruence).

Including a range of ages in our study allowed for investigation into age-related changes in children’s use of SLSs. Previous research has tended to focus on a behaviour of interest in more limited age groups (e.g., only including 5-year-olds: McGuigan, 2013; Wood et al., 2012), therefore precluding investigation into developmental trajectories or age-related changes. Our results suggest that the transition from implicit SLSs to explicitly metacognitive SLSs begins at around 6 years. Compared to 7- and 8-year-olds, the proportion of information seeking errors made by 6-year-olds was still relatively high. However, for those 6-year-olds who did select the target video, their appropriate information use was at a level comparable to that of 7- and 8-year-olds for both demonstration types. This suggests that 6-year-olds who selected the target video were likely to have been motivated by a reasoned understanding of the relevance of the information or were at least able to recognise its value once the information had been received, indicative of explicitly metacognitive SLSs. Indeed, the information seeking error rate in 6-year-olds (see Figure 2) could be a reflection of the transition from implicit to metacognitive SLSs, with some children having not yet developed the capacity to use such strategies so relying on less cognitively demanding implicit SLSs. This interpretation would be consistent with reports of advances in metacognitive understanding in partial exposure tasks (Rohwer et al., 2012) and tasks that require children to ascribe knowledge to others (Ruffman & Olson, 1989) from 6 years. Thus, we propose that cognitive advances in explicit metacognition are implicated in the ability to recognise the value of social information. These advances are related to the ability to reflect on one's own knowledge and, importantly, the ability to recognise when more information is needed to solve a particular problem, what that information is, and from where or whom it can be acquired. However, as we used age as a proxy for cognitive development and developmental trajectories of other cognitive capacities overlap with that of explicit metacognition, we cannot say for certain which cognitive developments might be implicated in the emergence of the capacity for explicitly metacognitive SLSs.

Finally, by using the explicit reasoning questions, we examined children's own understanding of the strategies they were employing when approaching the task. The results not only show that explicit reasoning improves with age, but also that higher rated reasoning responses were associated with better performance in both the information seeking and information use aspects of the task. In both cases we found that older children gave more responses that reasonably justified the selections they had made. This supports our assumption that appropriate information seeking (observed in older children) is indicative of employment of an explicitly metacognitive strategy rather than reliance on implicit biases (observed in younger children) which are not related to an understanding of the potential value of information. As alluded to in the introduction, implicit SLSs may sometimes be explicitly represented. As such, information seeking responses that were rated as reasoned but incorrect provide an interesting insight into children's interpretations of their model selections. This type of response revealed explicitly represented justifications for particular model selections that were not linked to the potential to provide valuable information. Rather, these responses indicated use of strategies related to alternative superficial cues, that, for example, fit with the model-based biases we observed in the results (e.g., response from 5-year-old female: *'they were all girls'*). While these responses are explicitly represented, they are not explicitly metacognitive. Thus, these results are consistent with the account we have presented proposing that younger children relied on implicit SLSs that were not driven by reasoning about the value of the available information.

Overall, these findings suggest that children, like adults, can actively seek out relevant information in order to identify appropriate social models. Here we presented evidence that 3- to 8-year-old children's capacity to actively seek out appropriate social information from multiple potential models exceeds what we would expect to see if they were selecting a model at random. We believe that our results suggest that older and younger children might have been using different SLSs in the information seeking phase of the task. In older children (7- and 8-year-olds), social information seeking appears to have been driven by explicitly metacognitive SLSs pertaining to an understanding of the information required and the potential value of the information that could be provided. By contrast, younger children appeared to rely on less cognitively demanding heuristic model-based biases (implicit SLSs) due to having not yet developed the capacity to use reasoning-based strategies. The pattern of results observed in this study highlights the need for and the benefits of expanding developmental social learning paradigms to take account of information seeking as well as information use. By doing so we provided evidence of an age-related transition from use of implicit to explicitly metacognitive SLSs.

Finding that older, but not younger, children were making reasoned choices, based on the value to themselves of social information, supports Heyes's (2016c) proposal that explicitly metacognitive SLSs are experience-dependent and develop relatively late in children. The relatively late emergence of the capacity to use reasoning-based, or explicitly metacognitive, SLSs suggests that this may be a cognitive mechanism that is unique to humans and is unlikely to be observed in animals. Thus, we propose that our findings are consistent with the interpretation that reasoned understanding of social information could account for distinctively human cumulative culture (Baldwin & Moses, 1996; Dunstone & Caldwell, 2018; Heyes, 2016c). Real world instances of the evolution of cumulative culture are likely to involve learners acquiring social information that is very limited in its availability, almost by definition. Thus, it is not unreasonable to suggest that capacities for appropriate information seeking could determine whether or not a particular population exhibits cumulative culture. The late development of this skill, as identified in our findings, suggests that appropriate social information seeking may have been overlooked as a significant cognitive challenge to fully benefit from others' knowledge. It is possible therefore that this could provide part of the explanation for the phylogenetic distribution of cumulative culture.

Chapter 4:

Age-Related Transition From Simple Associations to Reasoning-Based Social Learning Strategies in 4- to 8-Year-Old Children

---

The results of the study presented in chapter three indicated that children's appropriate information seeking appeared to be driven by a reasoned understanding of the information required to achieve their goal only from 7 to 8 years old. Therefore, older children demonstrated understanding of the value of others as sources of information and used explicitly metacognitive SLSs to appropriately seek out relevant information. By contrast, younger children seemed to rely on implicit SLSs, in this case heuristic biases for gender congruent informants.

However, the task employed in chapter three did not require participants to reason about the mental states of others. Instead, it was concerned with whether children were able to reason about the information that was required in order to solve a particular problem and whether they could determine who had the potential to provide relevant and valuable information. Thus, the study reported in chapter four aims to address whether children can use explicitly metacognitive SLSs in cases in which it is necessary to consider others' mental states. In particular, chapter four investigates the developmental trajectory of children's ability to identify potential informants' suitability as a source of information based on an understanding of the informant's knowledge, in this case privileged perceptual access.

The following chapter has been submitted for publication to the Journal of Comparative Psychology and is currently under revision for resubmission, the reference is given below. The chapter is presented in its originally submitted form. Data and analysis code are available in the OSF repository: [https://osf.io/et6ub/].

Blakey, K.H., Renner, E., Atkinson, M., Rafetseder, E., and Caldwell, C.A. (Under Revision). Age-related transition from simple associations to reasoning-based social learning strategies in 4- to 8-year-old children (*Homo sapiens*).

Abstract

To differentiate the use of simple associations from the use of explicitly reasoned selective social learning, we can look for age-related changes in children's behaviour that might signify a switch from one social learning strategy to the other. We presented 4- to 8-year-old children (N = 109) with a task in which they could at first succeed by forming an association or inferring a rule. However, following a switch in the scenario, success required explicit reasoning about the informants' potential to provide valuable social information based on their perceptual access to a critical event. The switch occurred once a proficiency criterion (five consecutive correct responses) was reached, or after 10 trials. Older children were more likely to reach criterion. Of all children meeting criterion, older participants were more likely to select a knowledgeable informant in the switch phase trials. This suggested that some of the younger children who had succeeded in the pre-switch trials had inferred rules or formed associations based on superficial, yet salient, visual cues, whereas older children had made the link between perceptual access and the potential to inform. This late development and apparent cognitive challenge are consistent with accounts proposing that such capacities may be linked to the distinctiveness of human cumulative culture.

Introduction

The ability to focus social learning on knowledgeable others has been proposed as a cognitive capacity underpinning distinctively human cumulative culture. This is because such abilities may enable explicitly metacognitive social learning strategies (Heyes, 2016c), which are proposed to fundamentally alter the pace and scope of cultural evolution (Dunstone & Caldwell, 2018; Heyes, 2018c). Social learning strategies (SLSs) – defined more broadly – have been considered responsible for much of the adaptive social learning behaviours found in humans and non-human animals (henceforth animals; Kendal et al., 2018; Rendell et al., 2011). However, it is generally accepted that human adults are good at identifying appropriate sources of social information and tracking ‘who knows’ (Heyes, 2016c; Nickerson, 1999), potentially revolutionising the strategic use of social information.

Despite interest in the idea that certain SLSs may be distinctively human, there have been few forays into the presence or emergence of these abilities in human children. The SLSs employed in all animal and much human behaviour are proposed to be based on implicit heuristics, including those learned through general-purpose associative learning processes. In contrast, explicitly metacognitive SLSs are claimed to be based on consciously represented reasoning-based rules. Determining whether individuals in particular studies were using associative learning processes or explicitly reasoned selective social learning has largely been applied in retrospect (Heyes, 2016a, 2017; Heyes & Pearce, 2015). However, as a means of potentially differentiating use of simple associations, or implicit heuristics, from use of explicitly reasoned selective social learning empirically, we can look for age-related changes in children’s behaviour that might signify a shift from reliance on one process to the other. Thus, this study investigates the ontogeny of explicitly metacognitive SLSs through exploring how children of different ages approach the task of identifying an appropriate person from whom to request information. Teasing apart the learning processes that underpin children’s selections of social sources could provide evidence to support the theory that reasoning-based explicitly metacognitive SLSs emerge relatively late in development. Tracking development may offer credence to proposals that such capacities are at least in part responsible for distinctively human cumulative culture.

Learning From the Right Others

Learning from others – social learning – can be very beneficial and is generally considered to be an adaptive capacity when it is present in a population. However, models have shown that adaptability is not improved simply by the inclusion of social learning. Rather, it is the ability to be selective and flexibly switch between social and individual learning that increases efficiency and adaptability (Enquist et al., 2007; Henrich & McElreath, 2003; Rogers, 1988). It follows then that specifically learning from appropriate others is key to successful social learning. To make best use of learning from others one should focus learning on those who possess relevant knowledge or experience. To some extent this targeted learning has been observed in humans and animals in the form of SLSs. These are flexible decision rules that influence when, what, and from whom social information should be acquired (Kendal et al., 2018; Laland, 2004; Rendell et al., 2011). Employing SLSs such as 'copy the majority', 'copy when uncertain', or 'copy the most successful' will generally improve the efficiency of social learning compared to learning indiscriminately from others (usually via blind copying) or learning by trial and error via individual learning. However, the majority of research has focused on the scope and functions of these SLSs, overlooking the cognitive mechanisms and the processes that underlie them.

There is widespread evidence of social learning (Hoppitt & Laland, 2013; Whiten, 2017b), use of SLSs (Horner et al., 2010; Jiménez & Mesoudi, 2019; Laland, 2004; Price et al., 2017), and culture (Allen, 2019; Aplin et al., 2015) in animals. In spite of this, there are marked differences between animals and humans. Specifically, there is disparity in regards to the capacity for cumulative culture (Caldwell et al., 2020; Dean et al., 2014). Cumulative culture refers to the accumulation of beneficial modifications to cultural traits over successive generations of learners, which results in increased functionality or efficiency (Dean et al., 2014; Mesoudi & Thornton, 2018; Tennie et al., 2009; Wilks & Blakey, 2018). One theory that has recently been proposed in an attempt to explain this seemingly distinctively human capacity is the idea that humans sometimes use explicitly metacognitive SLSs (Dunstone & Caldwell, 2018; Heyes, 2016c). Explicitly metacognitive SLSs are defined as consciously represented and reportable rules. This means that the agent is explicitly aware of the reasoning-based learning strategies that they are employing, including taking into account states of knowledge, ignorance, and uncertainty. These abilities may significantly influence how individuals seek out, attend to, and use social information, and as such may account for the distinctiveness of human cumulative culture (Blakey et al., 2020, Chapter 3 of this thesis; Dunstone & Caldwell, 2018). Learners who employ an explicitly metacognitive SLS have an appreciation of why their strategy is successful because they understand the causal link between another's knowledge or experience and their value as a source of information. That is, they have the capacity to recognise the potential value to themselves of others' knowledge based on an understanding of what makes that person a good source of information. Therefore, explicitly metacognitive SLSs are thought to enable tracking of *when* others are likely to have superior knowledge, and *who* of the available others is likely to be the best source of knowledge, so that social learning can be most efficient (Heyes, 2016a). As humans appear to be the only animals that focus their social learning by asking 'who knows', this targeted learning towards appropriate social sources is believed to, at least in part, facilitate cumulative culture (Heyes, 2018b; Shea et al., 2014).

Reasoning-based explicitly metacognitive SLSs form half of a dual-process account of selective social learning that suggests that there are distinct categories of SLSs, each based on specific types of decision rules (Heyes, 2016c). The other half of this account suggests that the majority of the behaviours that conform to SLSs in both humans and animals arise as a consequence of general-purpose associative learning processes or biologically selected biases. On the basis of relatively crude and imperfect heuristics, these *implicit SLSs* (as we will refer to them here) function to direct learning towards objects, agents, and events that are most likely to provide useful information. In contrast to explicitly metacognitive SLSs, they are not driven by a causal understanding of the potential value of social information, thus in principle, any salient association can be formed, or rule can be inferred. Though we refer to these strategies as implicit, it is possible that participants may sometimes infer and represent explicit rules that are not causally linked to the informants' potential to provide valuable information, and rather, are based on alternative more salient but relatively superficial cues. Thus, these rules may be explicitly represented; however, we suggest that they are not explicitly metacognitive due to the absence of a causal understanding of the relevance of the informants' mental states. Though we stated above that humans have the capacity to use explicitly metacognitive SLSs, it is likely that much of human behaviour is based upon these less cognitively demanding implicit heuristics. However, it is argued that while both humans and animals use implicit SLSs, only humans have the capacity for explicitly metacognitive SLSs.

Adult humans are - justifiably - assumed to be able to use explicitly metacognitive SLSs when seeking out and using social information (Caldwell et al., 2018; Vélez & Gweon, 2019). Therefore, it is not unreasonable to question the presence or emergence of these abilities in human children. Explicitly metacognitive SLSs are likely to be experience-dependent and learned through social interaction (Heyes, 2016c). As such they are likely to emerge relatively late in development after children have had the opportunity to learn through experience. Therefore, we would not expect to see explicit understanding of the value of social information, and how one might benefit from it, in infants or very young children. Given the assumption that adults do have this capacity, it follows that this understanding develops during the course of childhood. It may be possible to differentiate the SLSs that children are using, and identify age-related changes in behaviour, that might signify a switch from implicit to explicitly metacognitive strategies. This would provide insights into the cognitive capacities upon which these abilities depend, and therefore might also shed light on the reasons for the apparent absence of these abilities in animals.

The Learning Processes That Underpin Who Children Learn From

SLSs in the broadly-defined sense (also referred to in the literature as ‘learning heuristics’, or ‘transmission biases’) have been extensively researched in young children, especially in relation to ‘who’ strategies (see reviews by Price et al., 2017; Wood et al., 2012). These studies have tended to focus on children’s responses to social information (usually their propensity to copy, and their preferences for what and whom to copy) that has been passively acquired, rather than information that they have actively sought out. Similarly, literature on the phenomenon of ‘selective trust’ entails children choosing between models on the basis of social and/or epistemic characteristics (Koenig et al., 2004; Sobel & Kushnir, 2013) for which they have a pre-existing bias (from personal experience or phylogenetic history) or a prior history of associations (established as part of the experimental paradigm). While useful in terms of understanding how children use social information and who they learn from, in both lines of study (SLSs and selective trust) children's behaviour is consistent with a number of different potential mechanistic explanations, which include implicit associative learning and biologically selected biases, as well as explicitly reasoned metacognitive strategies. Therefore, they offer little scope for determining which of these is responsible. Moreover, Heyes (2017) proposed that the empirical evidence of SLSs in young children (under 4 or 5 years old) can largely be explained as a result of associative learning processes. We propose that only by setting the different potential mechanisms in conflict with one another can we really discover which is responsible. For example, to distinguish between implicit associations or biases and explicitly metacognitive reasoning, we could create a situation in which each potential mechanism will result in a different preference.

To date there have been very few studies that have specifically set out to track the development of explicitly metacognitive SLSs. However, recent evidence does indicate an age-related transition from the use of implicit heuristics to explicitly metacognitive SLSs that, as expected, occurs relatively late in childhood (Blakey et al., 2020, Chapter 3). Blakey et al. (2020, Chapter 3) examined children's understanding of what information they themselves required to solve a particular problem. This was achieved by examining children's model selections to see whether they were based on an understanding of the potential value of the information the model could provide, or on superficial, yet salient, characteristics such as age or gender. The results showed that, only at 7 to 8 years old did children's responses appear to be driven by reasoned understanding of the information required to identify an appropriate social model, consistent with an ability to use explicitly metacognitive SLSs. Younger children (3 to 6 years old) instead relied upon imperfect heuristic biases (in this case model-based biases for gender congruence), likely due to having not yet developed the capacity to use explicitly metacognitive strategies. However, that task did not require children to reason about the mental states of the potential informants. Rather, it focussed on children's recognition of the information they needed to complete a particular problem and their understanding about who could provide valuable information about that problem (i.e., faced with the same, rather than an alternative, problem). Thus, in the current study we sought to examine whether children take into account reasoning about others' mental states when judging 'who knows'. In particular, we wanted to investigate the development of the ability to identify an appropriate source of information based on a causal understanding of the link between the informants' knowledge or experience and their value as a source of information (i.e., tracking who knows). To test this, we created a situation in which different strategies for success could be learned, or inferred, across iterated trials during a training phase. We then examined what had been learned by inducing a switch in the scenario, intended to reveal the underlying rule or bias being applied by the participant. In these 'switch trials', a response informed by causal interpretations based on mentalistic reasoning about others' knowledge would directly conflict with a learned non-mentalistic association, or inferred rule, that had previously predicted success. Continued success following the switch would suggest that responses were informed by explicitly metacognitive mentalistic reasoning about others' knowledge, while lower success would indicate perseverative use of a learned association, or inferred rule, based on superficial cues.

While implicit SLSs may involve forming an association, or inferring a rule, between actions or objects that are predictive of success (Haselgrove, 2016), explicitly metacognitive SLSs require reasoned causal understanding as well as appropriate application of the rule (Caldwell, 2018; Dunstone & Caldwell, 2018). Reasoned understanding related to identifying relevant sources of social information can also be considered as reasoning-based information seeking. Similarly to Heyes (2016c), Baldwin and Moses (1996) proposed that motivations for seeking out appropriate social models are likely to be based on advanced metacognitive abilities. They highlight that to seek out social information effectively and appropriately a seeker must recognise the information that is needed, and from where, or whom, that information can be acquired. Of course, there are a number of prerequisites to recognising sources of social information. Seekers need to understand that others are sources of information and can thus provide information about particular objects or events, and importantly they should be able to appreciate this prior to the information being offered. It is also necessary for a seeker to understand concepts such as knowledge formation and ignorance in oneself as well as in others, including being aware of the potential for differential knowledge states. Dunstone and Caldwell (2018) recently postulated that explicit understanding of others' mental states may offer more flexibility with regards to ascertaining who the best informant is likely to be, posing a significant advantage in social information use. The abilities and understanding highlighted here are most likely learned, emerging with increasing experience of social interaction. The suggestion that these abilities are necessary precursors to reasoning-based information seeking fits with the expectation that this capacity, and thus explicitly metacognitive SLSs, becomes available to children relatively late in development. We discuss the developmental trajectory of these sociocognitive and metacognitive abilities in the following section.

Cognitive Requisites for Explicitly Metacognitive SLSs

Developing Understanding of Knowledge and Ignorance

From around 3 years old children begin to show evidence of understanding concepts of their own knowledge formation and ignorance; however, it is not reliable. Though they have the ability to evaluate their own knowledge and ignorance when they have full or no access to relevant information (Pratt & Bryant, 1990; Rohwer et al., 2012), issues with reliability arise when considering *how much* they know when they have access to only some of the information. In both verbal (Sodian & Wimmer, 1987; Wimmer et al., 1988) and behavioural (Kloo & Sodian, 2017; Rohwer et al., 2012) partial exposure tasks (participants are exposed to a range of objects, but then cannot see which of the objects is being put inside a box) children under 6 years old consistently overestimate their own knowledge. One potential explanation for younger children's poor performance compared to total ignorance tasks is that children under 6 years old lack an understanding of how knowledge is formed. Without the recognition that knowledge is formed on the basis of access to information (Kloo & Rohwer, 2012; Rohwer et al., 2012), children appear to rely on a 'sense of knowing' when evaluating their own ignorance. However, children reflect an understanding of the relevance of access to information in knowledge formation in their behaviour earlier (from 3 to 4 years old) than they do in their explicit judgements of the source of the information (from 5 years old; Robinson et al., 2008; Sodian et al., 2006). Similarly, behavioural tests of implicit metacognitive ability suggest that although infants (Goupil et al., 2016), young children (Bernard et al., 2015), and animals (Call & Carpenter, 2001; Neldner et al., 2015) react to the state of ignorance, they do so without necessarily recognising a metacognitive awareness of that state. This ability appears to precede the ability to explicitly evaluate knowledge states, and these different developmental timelines correspond with the idea that explicit responses are underpinned by different processes that come online later in development (explicitly metacognitive SLSs). Thus, while younger children's responses may sometimes give the appearance of an understanding of states of knowledge, their struggle to report these states suggests otherwise. Rather, their responses likely rest on less cognitively demanding processes (e.g., implicit SLSs, in the case of adaptive biases in their social learning).

Understanding the Role of Perceptual Access in Knowledge Formation

Appreciating the value of information that another individual can provide requires an understanding of how knowledge is acquired both by oneself and by others; this includes recognising the link between access to information (commonly and saliently afforded via visual or auditory perception) and knowledge formation. Thus, to evaluate whether young children appreciate how knowledge is acquired, it is necessary to consider the age by which children recognise that perception leads to knowledge formation. Such understanding is likely to be an important precursor to determining other individuals' value as sources of information. Identifying whether children (from around 3 to 4 years old) recognise that knowledge is formed via access to information allows us to ascertain whether their behaviour demonstrates, as is claimed, that they can differentiate knowledge and ignorance.

Here we focus on the role of visual perceptual access in knowledge formation as this is arguably the most salient to young children, and most relevant to the procedure used in the current study. While there is evidence that children as young as 2 years old show sensitivity to others' knowledge states as a result of their visual experience (Dunham et al., 2000; O'Neill, 1996), the majority of research suggests that from 3 to 4 years old children can use others' perceptual experience to determine their knowledge or ignorance (Perner, 1991; Pillow, 1989; Pillow & Weed, 1997; Pratt & Bryant, 1990; Sodian et al., 2006). However, Ruffman and Olson (1989) found that while 3- to 4-year-olds were able to assess others' perceptual access correctly, only from 6 years old were children able to report another's knowledge. O'Neill et al. (1992) offered a similar finding, indicating that 3-year-old children have some understanding that knowledge and perceptual access are associated, but not until around 4 or 5 years old do they develop an understanding of how perception and knowledge are causally related (i.e., seeing leads to knowing). These results suggest that understanding perceptual access precedes the ability to attribute knowledge. The conflicting evidence regarding the age at which children develop an appreciation of the causal influence of visual perception on knowledge formation suggests that such recognition develops gradually until around 5 years old and thus is likely to be experience-dependent, at least in part. If younger children are less likely to understand the fundamental role that perceptual access plays in knowledge formation, it suggests that the behaviours they display that appear to demonstrate such understanding may be based on different learning processes (possibly learned associations or heuristics) than those required to be able to attribute knowledge. It follows that attributing knowledge on the basis of perceptual access may require explicit metacognitive reasoning which appears to develop relatively late.

Determining Others' Value as Sources of Information

It is also vital for understanding the development of reasoning-based social information seeking to establish whether children can actually *use* their understanding of others' knowledge states, as inferred from perceptual access, to ascertain others' value as a source of information. A variety of studies report that children from 3 to 4 years old preferentially select or report agents as more knowledgeable or trustworthy on the basis of their perceptual access (e.g., Brosseau-Liard & Birch, 2011; Fedra & Schmidt, 2019; Koenig et al., 2019; Mills, 2013; Terrier et al., 2016). Yet research investigating whether children can use understanding of others' knowledge states to determine their value as a source of information is scarce. An exception is Robinson et al. (2011), who explored whether young children recognise that they can gain valuable knowledge from others who have had perceptual access to something that they have not. Primary results revealed that 3- and 4-year-old children behaved as if they understood that knowledge could be gained from an informant when the informant had had an experience (looking inside a box) that the children themselves had not. In follow-up tasks children were required to choose who-to-ask or whether-to-ask one of two puppets (knowledgeable vs ignorant) to provide desirable information. In the task, children observed the experimenter allow one puppet to look inside a target box, and the child showed the contents of a pencil case to the other puppet. After the demonstration, the experimenter told the child that they had to say what was inside the box. In the who-to-ask condition children had to choose one of the puppets to help them, while in the whether-to-ask condition they could choose one of the puppets or make their own guess. The results revealed that 5;3- to 6;2-year-olds asked the knowledgeable puppet more frequently than chance in both conditions. In contrast, even when they were forced to ask one of the puppets for help, in the who-to-ask condition, younger children (4;3- to 5;2-year-olds) were no better than chance at asking the puppet with the relevant experience, indicating that the younger children failed to recognise the *value* of the others' knowledge despite being able to report the puppets' knowledge based on their perceptual experience. Similarly, O'Neill et al. (1992) also reported that 4- and 5-year-old children found assessing knowledge states relative to their perceptual access more challenging when they were required to compare two puppets' knowledge states rather than evaluate the knowledge state of a single puppet. From the evidence presented here it appears that while children under 5 years old are able to identify others' perceptual access, they may struggle to recognise that perceptual access grants privileged knowledge to an individual such that they become a valuable source of information. The cognitive challenge posed by these metacognitive capacities may preclude younger children from identifying the most appropriate source of social information. Hence, we suspect that while older children are able use others' perceptual access to explicitly reason their knowledge state, younger children may instead be relying upon implicit SLSs.

Differentiating Implicit and Explicitly Metacognitive SLSs Empirically

Returning to our goal of investigating when children develop the ability to shift from using implicit SLSs to flexibly employing explicitly metacognitive ones, in search of a suitable study design we considered the efficacy of existing paradigms. When experiencing situations of ignorance or uncertainty, individuals who understand concepts such as the potential to acquire desired information or knowledge from others are more likely to request information from others who they believe to be better informed (Baldwin & Moses, 1996). Yet, if the goal is to determine whether individuals are using explicitly metacognitive SLSs to determine their choice of informant, it is necessary to provide a scenario in which the appropriate choice cannot be predicted by alternative environmental cues that do not require consideration of others' knowledge states.

Several studies have used 'perspective-taking' paradigms to investigate whether non-human primates (Povinelli et al., 1990, 1991) or young children (Povinelli &

DeBlois, 1992) understand the relationship between seeing and knowing. Results revealed that chimpanzees and 4-year-old children, but not rhesus macaques or 3-year-old children, responded preferentially to information provided by informants that were knowledgeable as opposed to ignorant. Though argued to have provided evidence that non-human primates and young children are able to attribute knowledge and ignorance to others, Heyes (1998) contested the validity of such conclusions. Heyes claimed that preference for the knowledgeable informant could be explained by formation of an associative rule across trials, and thus proposed a novel methodology to address such issues.

Heyes's proposed method involved giving an individual first-hand experience of a novel apparatus (barrier) that granted or denied perceptual access. Though seemingly the same from a distance (other than being distinct colours), one barrier would be transparent (permitting seeing) while the other would be opaque (preventing seeing). Following the self-experience phase, two agents, each assigned one of the barriers, would attend to a reward hiding process after which the individual's choice of agent (seeing or not-seeing) would be recorded. When the perceptual access of others is not visually salient, such as when it is granted or denied according to the opacity of a barrier (e.g., goggles/blindfolds/screens), participants must use personal experience of the properties of the barriers to identify the perceptual access it affords to others (Lurz & Krachun, 2019). Accordingly, preference for the agent who can see through their barrier would suggest attribution of knowledge states on the basis of self-experience of the barrier's properties, as opposed to a response that had been reinforced by prior experience of similar situations (as might be true for a preference for e.g., head orientation).

Using Self-Experience to Infer Others' Perceptual Experience

Several studies have since employed methods inspired by the paradigm set out by Heyes (1998; and adapted by Povinelli & Vonk, 2003, 2004) to investigate whether infants (Meltzoff & Brooks, 2008; Senju et al., 2011; Teufel et al., 2013) or great apes (Kano et al., 2019; Karg et al., 2015) can use their own perceptual experience to infer the perceptual experience of others, and so too, what others know. These studies have concluded that both infants (from 18 months old) and great apes can use their own perceptual experience to determine others' perceptual access. However, arguably these studies capture implicit metacognitive abilities, since they rely upon anticipatory-looking paradigms or behavioural reactions, which, as previously discussed, appear to precede explicit reasoning regarding others' knowledge states. Importantly, Teufel et al. (2013) highlighted that although 2-year-olds were able to use their personal experience to infer others' visual perception, they appeared to lack the causal understanding of the link between visual perception and knowledge formation. This discrepancy was evident from children's indiscriminate use of pointing to request a hidden object from knowledgeable and ignorant parents (knowledge was afforded/denied to parents via transparent/opaque glasses during the hiding event). These results suggest that children of this age have not yet developed the ability to attribute knowledge to others using their own past perceptual experience to infer the perceptual access of others. In comparison, adults have the capacity to use reasoning-based strategies. For instance, if adults are ignorant about a critical event they can use their past experience of similar events to infer the experience (perceptual access) of potential informants and attribute knowledge to appropriate individuals (Teufel et al., 2009).

Despite the wealth of research discussed above, little consideration appears to have been given to the question of whether children can use their own perceptual experience, not only to identify others' knowledge or ignorance, but also to ascertain and utilise the value of another individual as a source of information. Thus, the purpose of this study was to address this gap and investigate age-related changes in children's ability to use their own experience as a means to reason the perceptual access of potential informants to a critical event to which the children did not themselves have access.

To summarise, understanding one's own and others' perceptual access, and the ability to attribute knowledge to others, are necessary for understanding that perceptual access to a critical event grants privileged knowledge regarding the outcome of that event. If one's own perceptual access to a critical event was obstructed, we postulated that explicit reasoning may be required to identify an appropriate individual from whom to request information. While there is evidence, from the SLSs and selective trust literature, that very young children (Jaswal & Neely, 2006; Koenig & Harris, 2005; McGuigan, 2013; Rakoczy et al., 2010; Tong et al., 2020) and animals (Horner et al., 2010; Kendal et al., 2015; Ottoni et al., 2005) do sometimes select knowledgeable individuals, we propose that these selections are unlikely to be driven by an understanding of the causal link between another's knowledge and their value as a source of information. For example, biases for older, high-status, or reliable models may be the result of repeated exposure to the successes of models with these characteristics, thus resulting in learned associations or rule-like strategies (implicit SLSs). In contrast, adults employ strategies that enable them to reason about others' mental states when judging 'who knows'. Consequently, they appropriately attribute and request knowledge from others who have had access to desirable information. The flexibility with which adults engage this ability suggests that they reason the perceptual access of others and employ their understanding that perceptual access leads to knowledge to assess others' suitability as sources of information (explicitly metacognitive SLSs). Though implicit SLSs are necessarily broadly adaptive, being able to explicitly reason who is an appropriate informant is likely more advantageous due to its flexible nature and that it can be applied across contexts.

The Current Study

The aim of the current study was to investigate age-related changes that might reflect a switch from the use of implicit SLSs to the use of reasoning-based explicitly metacognitive SLSs. To do this we designed a task in which participants could at first succeed by forming a simple association, or inferring a rule, to predict success. The key element of this study was that, following a switch in the scenario, in order to continue to be successful participants were required to explicitly reason the informants' perceptual access to a critical hiding event. Following the switch, continued use of an association, or inferred rule, would result in imperfect responses, leading to lower success. Importantly, we specifically designed this study in such a way that success in the pre-switch trials was not limited by the capacity for mentalistic causal understanding. We wanted to provide participants with the possibility of being successful by learning across multiple trials, either by forming an association or inferring a rule, without the need for personal insight into such processes (i.e., through implicit SLSs).

To provide participants this opportunity to form an association, or infer a rule, we adopted an experimental approach more commonly used in studies of non-human primates. This involved exposing participants to repeated trials until they reached a proficiency criterion level (or completed the maximum number of trials), at which point they progressed to the switch phase of the task. Reaching criterion was taken to mean that children had either formed an association, inferred a rule, or that they had understood the causal relationship between the informants' perceptual access and their knowledge and were using this to identify a knowledgeable informant. We did not expect all children to reach criterion due to the relatively low number of trials children had in which to form an association, or infer a rule, compared to the number of trials and sessions afforded to other species using tasks similar in design. The switch was intended to generate a conflict between the response favoured by associations, or inferred rules, that predicted success in the pre-switch phase and the response favoured by an understanding of why another's behaviour is informative. Age-related changes in children's responses following the switch were expected to offer an insight into the learning processes being employed (i.e., implicit or explicitly metacognitive SLSs).

We were primarily interested in what the children who reached criterion were responding to in the switch trials, as whilst the appropriate causal mentalistic interpretation would continue to be successful in the switch trials, the other cues would no longer predict success. Additionally, because of what we know about children's developing capacity for appreciation of others' mental states, we expected to find an age-related difference in the SLSs being employed. We expected that older children would be more likely to use reasoning-based judgements to attribute value to the informants' knowledge on the basis of their perceptual access to the critical event (i.e., explicitly metacognitive SLSs). By contrast, we expected that younger children would struggle to employ such reasoning-based strategies, and therefore would be more likely to rely on simple associations, or inferred rules, learned over the course of the pre-switch trials (i.e., implicit SLSs). If children were using explicitly metacognitive SLSs, we would expect them to reach criterion in the pre-switch trials and continue to be successful in the switch trials due to a mentalistic causal understanding of the informants' potential to provide valuable information. If, however, children were instead relying on learned associations, or inferred rules, we would expect them to struggle to continue their success in the switch trials (despite reaching criterion) due to their success in the pre-switch trials not being driven by a causal understanding of the informants' mental states.

It was anticipated that this task would provide an insight into the cognitive mechanisms underlying children's use of social information and children's approach to learning from social sources. Previous research into the development of reasoning-based learning strategies suggests that such capacities are cognitively challenging and develop relatively late (Blakey et al., 2020, Chapter 3). Thus, if we were to find a similar developmental trajectory this may serve as evidence that reasoning-based learning strategies influence distinctively human cumulative culture. More specifically, it would be consistent with the interpretation that being able to fully benefit from others' knowledge is dependent, at least in part, on the ability to assess the suitability of others as sources of knowledge based upon their perceptual access to the desired information.

Method

Participants

The final participant sample comprised 109 children aged four to eight years (55 females; Mean age = 76.8 months, *SD* = 16.9 months, range = 48 to 107 months). Participants were visitors at the Royal Zoological Society of Scotland's Edinburgh Zoo. The nationality of the sample, as identified by parents/legal guardians, was predominantly British (> 90%). An additional 13 children were tested but later excluded from analyses. These exclusions were due to researcher error (*n* = 3), missing data including non-completion of the task (*n* = 7), and task interference, including intervention, or distraction, by parents or other children due to testing taking place in a public area (*n* = 3). This study was granted ethical approval by the University of Stirling General University Ethics Panel (GUEP673), and informed written consent was provided by each participant's parent/legal guardian prior to data collection.

Materials

Two adult female informants followed a counterbalanced and randomised set of actions presented to them via a program written in PsychoPy v1.84.2 (Peirce et al., 2019) which was run on a Microsoft Surface tablet. Each adult was referred to according to the pattern of the t-shirt that they were wearing (one black with white polka dots, '*Spots*', and one with black and white stripes, '*Stripes*'). Two black wooden boxes (5.8 x 5.8 x 3.5cm) served as potential reward locations; a target box contained a reward object, and a non-target box was empty. The target box was identifiable to the adult informants via a discrete sticker on the back, unnoticeable to participants. Square Lego Duplo blocks (red, blue, green, and yellow) were used as the reward objects and a Lego Duplo baseboard on which participants could build a Lego tower was provided. Two types of free-standing wooden frames consisting of yellow and blue frames (35 x 35cm) and black opaque or black semi-opaque inner screens (30 x 30cm) were used to manipulate informants' perceptual access (see Figure 1; based on materials used by Karg et al., 2015). The blue frame with the semi-opaque screen permitted perceptual access when viewed straight on and at a close distance (this was the view children had in the experience phase; see Figure 1A), though when viewed at an angle or from a distance it appeared to be opaque (this was the view that children had in the pre-switch and switch phases; see Figure 1B). The yellow frame with the opaque screen prevented perceptual access and was used to occlude an informant's view. Another larger yellow frame (45 x 35cm) fitted with a black opaque inner screen (40 x 30cm) was used to occlude the participant's view. A laminated card (one purple, one green) was fixed to each side of the testing table so that the colour of the sides of the table could be used as a potentially salient visual cue in the pre-switch phase of the task.

Figure 1

Illustration of the Perceptual Access Afforded by the Opaque (Yellow) and Semi-Opaque (Blue) Frames When Interrupting the View of Another Object

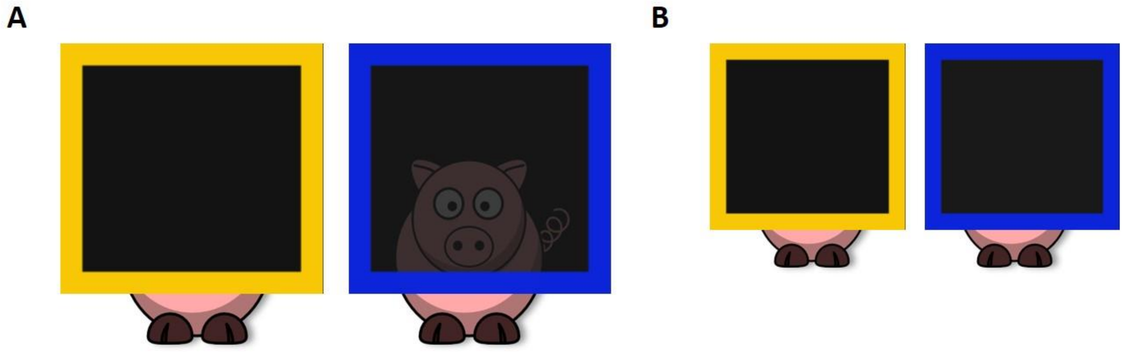

Note. A: when viewed close up, illustrative of child's view during the experience phase; B: when viewed at a distance, illustrative of child's view during pre-switch and switch phases.

Design

Participants took part individually in a single testing session for which they received a sticker. Due to testing taking place in public area, a familiar adult was present but was asked not to interact with the participant during the task. The experimenter sat beside the participant, and the two informants sat next to each other on the opposite side of the table, facing the experimenter and the participant. Ideally, the informants would have sat at the corners of the table so that the participant would view the frames at an angle. However, practical constraints such as the requirement for the Spots and Stripes to repeatedly swap places and that they both had to see the instructions on the tablet, meant that they had to sit close together and opposite the participant. In spite of this compromise, the distance at which they sat ensured that both the opaque and semi-opaque frames appeared to be opaque (see Figure 1B).

The task comprised three sequential phases: an experience phase, a pre-switch phase, and a switch phase (see Appendix C for a verbal script). The experience phase was included to familiarise the participants with the properties of each of the different frame types thus providing information about the perceptual access afforded by each type of frame. The purpose of the pre-switch phase was to present a scenario in which participants could form an association, or infer a rule, between one of a number of potentially salient visual cues and the desired outcome which did not require explicit reasoning of the informants' perceptual access. Following the pre-switch phase there was a critical switch in the task scenario which rendered associations formed, or rules inferred, in the pre-switch trials ineffective. The trials in the switch phase were intended to generate a conflict between the response favoured by associations, or inferred rules, that predicted success in the pre-switch phase and the response favoured by an understanding of why another's behaviour is informative (i.e., explicit reasoning regarding the causal link between informants' perceptual access to the critical event and their potential to provide valuable information). Following completion of the switch phase, participants were asked to provide explicit verbal reasoning for their selections: the experimenter asked, "*How were you choosing who you wanted to tell you where the Lego was?*" Participants' responses to the reasoning question were categorised into four levels: no response, non-reasoned response, reasoned but incorrect response, and reasoned correct response.

Procedure

Experience Phase

The task began with participants being introduced to the two types of frames. The experimenter explained and demonstrated that "*You can see through frames that are blue, but you can't see through frames that are yellow*" (see Figure 1A). Participants were invited to try looking through both types of frame to see which one they could see the experimenter through. They were then asked to identify which frame they could see the experimenter through, and which frame the experimenter could see the participant through. Responses to these perceptual access questions were recorded on the tablet by one of the adult informants as a check of participants' understanding of their own and others' perceptual access. Any errors were corrected by the experimenter before moving on to the pre-switch phase.

Pre-Switch Phase

Following the experience phase participants were told that they were going to "*play a finding and building game with Spots and Stripes*", and the experimenter pointed at the adult informants in turn to identify them. At the start of each trial participants were shown two boxes, one which contained a Lego Duplo block (target) and one which was empty (non-target). The experimenter explained "*You and Spots and Stripes are going to try and find the Lego to build a tower*". The experimenter explained that before the Lego was hidden "*One frame will go in front of Spots, one frame will go in front of Stripes, and this frame will go in front of you.*" The blue (semi-opaque) frame was placed in front of the knowledgeable informant and the yellow (opaque) frame was placed in front of the ignorant informant; once both informants' frames were in place the second larger yellow (opaque) frame was placed in front of the participant which occluded their view of the table and the informants' frames. When all frames were in place the experimenter mixed around the open boxes in the centre of the table between the participant's frame and informants' frames before closing the boxes so that the Lego was hidden. After the Lego had been hidden, the frame in front of the participant was removed from the table and the informants moved the frames in front of them to the sides of the table (see Figure 2A for an example of a trial). The experimenter asked the participant "*Who do you think will tell you where the Lego is hidden?*" The informant chosen by the participant (either verbally or by pointing) was recorded on the tablet. We chose to ask participants to select an informant without any information being offered by the informants (e.g., via the informants pointing at the boxes) due to previous research finding that pointing interferes with children's ability to discriminate between informants (Palmquist et al., 2012). Regardless of whether the knowledgeable or the ignorant informant was chosen, the knowledgeable informant opened the target box to reveal the Lego inside (i.e., these trials were 'no-risk' as the reward was revealed regardless of which informant was selected). The participant was then able to retrieve the Lego and begin building a tower.

The same procedure was repeated for up to ten pre-switch trials, beginning with the participant being shown that one of the two boxes contained Lego and that the other was empty. The starting side of the knowledgeable informant (side of the semi-opaque frame) was randomly generated at the start of the pre-switch trials. Once set by the random generator, the sides of the knowledgeable and ignorant informants (i.e., the sides of the semi-opaque and opaque frames) remained constant throughout the pre-switch trials. The sides that Spots and Stripes sat on randomly alternated between trials, thus although the position of the knowledgeable informant (as indicated by the blue semi-opaque frame) was consistent, the identity of the informant (Spots or Stripes) was not. The positions of Spots and Stripes were pseudo-randomly determined for each trial, such that if the participant completed all ten pre-switch trials there would be an equal distribution of trials in which Spots and Stripes were on each side. The initial frame position and all variations of informant positions were indicated to the informants and the experimenter on the tablet which was not visible to the participant.

Figure 2

Schematic Illustration of the Task From the Perspective of the Participant, Including Placement of Frames and Boxes.

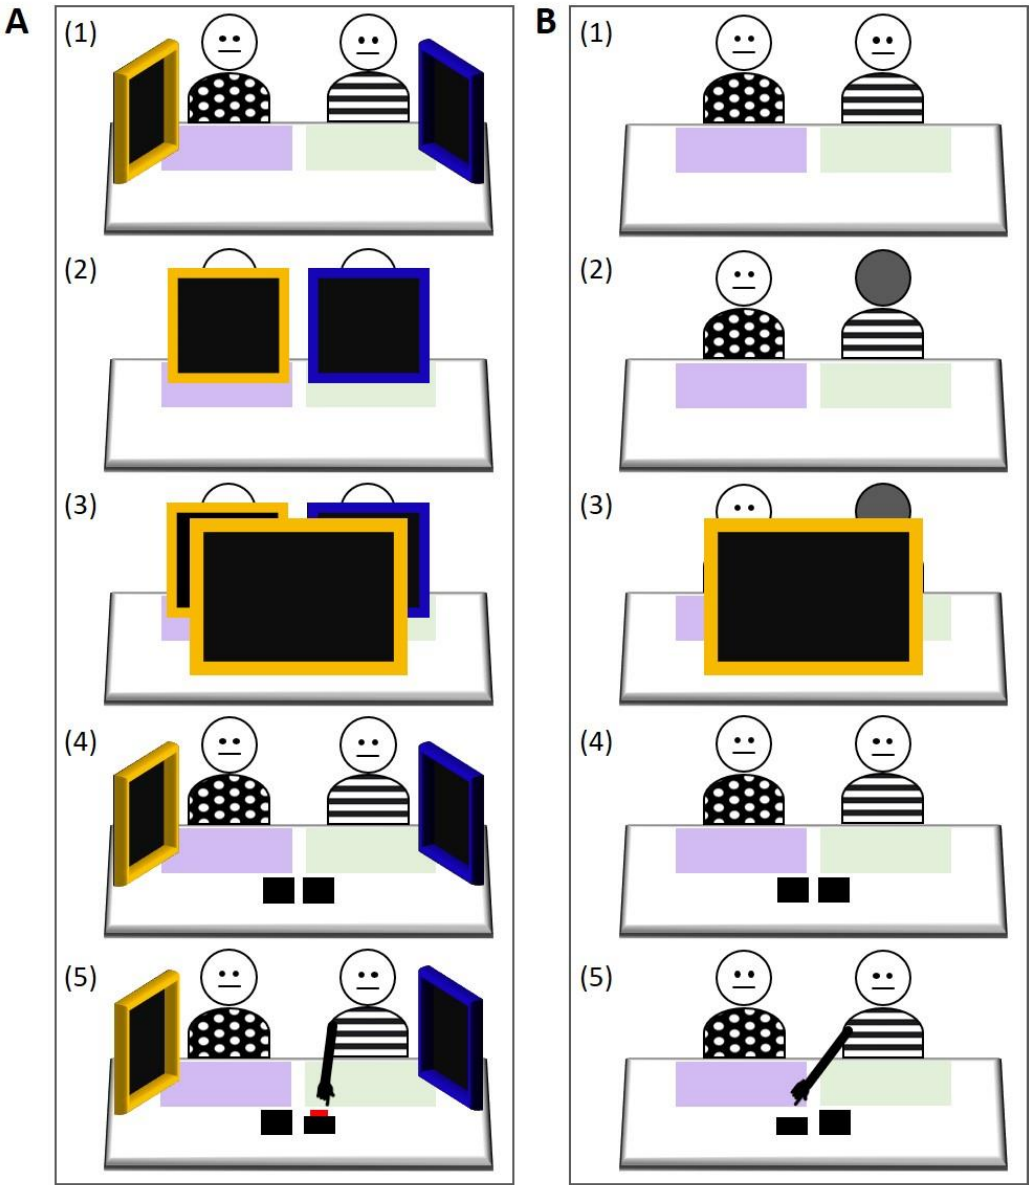

Note. A: example of trials in which perceptual access is indicated by yellow (opaque) and blue (semi-opaque) frames; B: example of trials in which perceptual access is indicated by the informants facing or not facing the critical event (i.e., the ignorant informant faced away). The stages illustrated are: (1) start of the trial, (2) A: frames placed in front of informants, B: ignorant informant turns away, (3) yellow (opaque) frame occludes participant's view, behind the frame two boxes (one containing Lego) are mixed around and closed, (4) A: frames removed, B: ignorant informant turns back around, (5) participant selects informant, reward revealed or not revealed depending on phase and selection.

Participants progressed to the switch phase of the task either after they had selected the knowledgeable informant in five consecutive pre-switch trials (met criterion), or after they had completed ten pre-switch trials. To reach criterion participants could have explicitly reasoned the informants' potential to provide valuable information based on their perceptual access to the critical event (when the Lego was being hidden), or they could have used superficial, yet salient, visual cues to form an association, or infer a rule, that predicted success. Appropriate associations or rules could have been based on the colour of the side of the table, the side of the table, or the colour of the frame. Any one of these cues may have been more salient to participants than perceptual access and in this phase would have had the same desired outcome. The 'no-risk' nature of this phase was employed to reduce the number of trials that may be required for participants to form an association, or infer a rule, that predicted success. A reasonably low number of trials, relative to the number of trials and sessions afforded to other species using similar paradigms, was considered preferable following piloting of this study (see Appendix D) due to the young age of the participants and the need for them to remain motivated to complete all phases of the task. It is important to note that not all children were expected to reach criterion due to the very limited number of trials in which to form an association, or infer a rule, that predicted success.

Switch Phase

Following completion of the pre-switch phase (either through reaching criterion or completing the maximum number of ten trials) there was a 'switch' in the task scenario. The switch referred to the random switching of the positions of the knowledgeable and ignorant informants at the beginning of each trial. The switch phase comprised five trials that followed a similar format to the pre-switch phase trials, however there were two key changes to the procedure.

The first change was that the positions of the informants (as indicated by the frames) no longer remained constant. As the positions of the informants had remained constant throughout the pre-switch phase, switching the informants' positions rendered most learned associations, or inferred rules, ineffective (with the exception of frame colour). Therefore, children were required to reason about the informants' perceptual access to continue to be successful. In the first trial of the switch phase the positions of the knowledgeable and ignorant informants always switched from their pre-switch phase positions. Thus, the positions of the knowledgeable and ignorant informants (as indicated by the frames) were the opposite of their pre-switch phase positions. In the following four trials, the side of the knowledgeable informant was randomly assigned, so sometimes the positions of the knowledgeable and ignorant informants switched and sometimes they did not.

As was the case in the pre-switch trials, for the first three trials in the switch phase, the informants' perceptual access was indicated by the two different frame types, therefore the positions of the yellow (opaque) and blue (semi-opaque) frames were switched randomly at the beginning of each trial. However, in the final two switch phase trials perceptual access was indicated by the knowledgeable informant facing, and the ignorant informant facing away from, the critical event (see Figure 2B for an example of a trial). The final two trials with no frames were included to explore whether participants were truly responding to the informants' perceptual access, or whether they had simply made the connection between the reward and the frame, or frame colour, that they could have been continuing to use as a predictive cue. Therefore, to be successful across all switch trials, participants were required to explicitly reason the informants' perceptual access to the critical event and understand why another's behaviour is informative.

The second key change to the procedure related to the selection that participants were asked to make. The experimenter signalled a change in the importance of the participant's decisions: "*This time you can choose who you want to tell you where they think the Lego is. Try hard to pick the right person, or you might not get the Lego this time.*" Instead of being asked who they thought would tell them where the Lego was hidden, the participant was asked "*Who do you* ***want*** *to tell you where they think the Lego is hidden?*" The change to this question was also reflected in the informants' reactions to participants' selections. In contrast to the pre-switch phase, the informant chosen by the participant selected a box and opened it to reveal its contents. If the knowledgeable informant was chosen then the target box was opened to reveal the Lego, while if the ignorant informant was selected the non-target box was opened and the Lego remained concealed.

Statistical Analysis

Analyses were performed using R (R Core Team, 2020), with generalised linear mixed effects analyses (GLMMs) performed using lme4 (Bates et al., 2014) with logit regression. P-values < .05 were accepted as statistically significant. The binary dependent variable in the analysis was the successful selection of the knowledgeable informant in the switch trials. Where specified as fixed effects the following variables were sum coded: met criterion (Criterion not met as −1, Criterion met as 1), the presence of the frames (No frames as −1, Frames as 1) and the number of pre-switch trials required to meet criterion (>5 trials as −1, 5 trials as 1). Age was centred and scaled to measure thousands of days. The random effects structure for each model aimed to include by-participant random slopes for all fixed effects and keep random effects structures 'maximal' where possible (following Barr et al., 2013). Where the 'maximal' model resulted in non-convergent or singular fit models, random slopes were removed, followed by random intercepts where necessary, until a convergent, non-singular model was obtained. Post hoc analysis was carried out using estimated marginal means using the *emmeans* package (Lenth et al., 2019). Post hoc results are given on the log odds ratio scale.

Results

The primary aim of the analysis was to assess the switch phase performance of children who met the proficiency criterion in the pre-switch phase of the task. This was to investigate whether children were able to use explicit reasoning about the informants' perceptual access to identify which informant was knowledgeable. We also hoped that we would be able to identify the SLSs that children were employing when approaching the trials in the pre-switch and switch phases. Successful performance was based on whether children selected the knowledgeable informant in a given trial. Overall, 52 children (48%) reached the proficiency criterion in the pre-switch phase (see Table 1). Reaching criterion was taken to mean that children had either formed an association, inferred a rule, or that they had explicitly reasoned the causal relationship between the informants' perceptual access and their knowledge, leading them to select the knowledgeable informant in five consecutive pre-switch trials. Continued success in the switch trials after reaching criterion could indicate reasoning of the informants' knowledge as a result of their perceptual access (i.e., explicitly metacognitive SLSs), while failure on switch trials after reaching criterion could suggest that children's success in the pre-switch trials was the result of learned associations, or inferred rules, not based on a causal understanding of the informants' knowledge (i.e., implicit SLSs).

Table 1

Number of Children Who Met Criterion by Age in Years (N = 109) and, for Those Who Met Criterion (n = 52), Whether Five or More Than Five Pre-Switch Trials Were Required to do so

		Criterion met		No. of pre-switch trials required to meet criterion	
	n	No	Yes	5 trials	> 5 trials
4 years	21	14	7 (33%)	3	4
5 years	23	13	10 (43%)	6	4
6 years	22	11	11 (50%)	8	3
7 years	23	11	12 (52%)	8	4
8 years	20	8	12 (60%)	9	3
All	109	57	52 (48%)	34	18

Switch Phase Performance

Effect of Reaching Proficiency Criterion

We first compared the switch trial performance of children who did and did not reach the proficiency criterion in the pre-switch trials. A GLMM was built for successful selection of the knowledgeable informant in each switch trial with fixed effects of age, met criterion, the presence of frames, and the interactions between these variables, and random intercepts of trial number and participant ID. A significant main effect of age ($b = 0.99$, $SE = 0.27$, $z = 3.61$, $p < .001$) indicated that older children selected the knowledgeable informant more often than younger children. A significant main effect of meeting criterion ($b = 0.75$, $SE = 0.14$, $z = 5.46$, $p < .001$) revealed that children who met criterion in the pre-switch phase selected the knowledgeable informant more often in the switch trials than children who did not meet criterion. There was a significant two-way interaction between age and meeting criterion (b = 0.58, SE = 0.27, z = 2.14, p = .033; see Figure 3). To clarify this interaction, we performed a post hoc analysis using *emmeans.* This revealed that the effect of meeting criterion was slightly more pronounced in older (upper quartile; b = 0.29, SE = 0.06, z = 5.15, p < .001) than younger children (lower quartile; b = 0.22, SE = 0.07, z = 3.30, p = .001). The GLMM did not provide evidence of an effect of the presence of frames (b = −0.08, SE = 0.13, z = −0.61, p = .544) on successful selection of the knowledgeable informant. This suggests that children were equally successful whether they were required to infer the informants' perceptual access using their experience of the properties of the frames or using the direction that the informants were facing.

Figure 3

Mean Number of Switch Trials in Which Children Selected the Knowledgeable Informant by Age in Years and Whether Children Met Criterion or Not

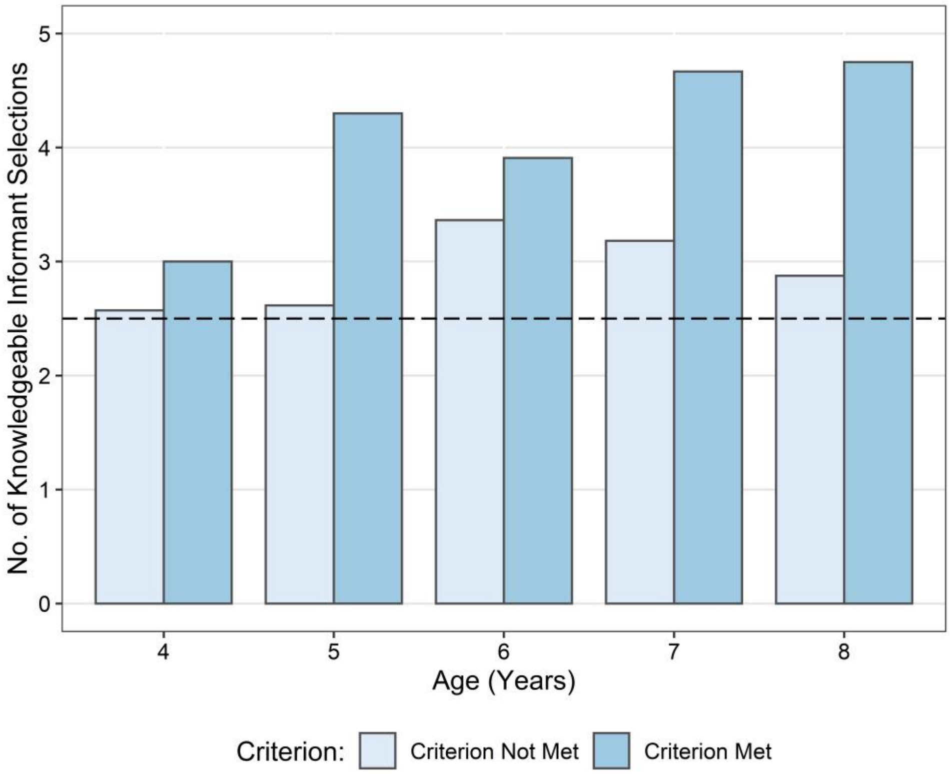

Note. Dashed line indicates chance performance.

For each participant we calculated the number of knowledgeable informant selections they made across the five switch trials. We then used these within-participant scores to compare performance to chance (50%) using one-sample *t*-tests. Children who met criterion selected the knowledgeable informant in an average of 4.23 switch trials (see Table 2); this was found to be significantly above chance $t(51) = 11.78$, $p < .001$, one-tailed. Children who did not meet criterion selected the knowledgeable informant in 2.89 switch trials, this was also found to be significantly above chance $t(56) = 2.36$, $p = .011$, one-tailed. Though both groups performed above chance, an independent-samples *t*-test showed that the difference between these two groups was significant $t(106.27) = 6.00$, $p < .001$, one-tailed.

Table 2

Mean Number of Knowledgeable Informant Selections Across the Five Switch Trials

		No. of knowledgeable informant selections	
	n	Mean	SD
Criterion not met	57	2.89	1.26
Criterion met	52	4.23	1.06
Criterion met in >5 pre-switch trials	18	3.56	1.38
Criterion met in 5 pre-switch trials	34	4.59	0.61

Age-Effects in Children Who Met the Proficiency Criterion

Having established that children who met criterion selected the knowledgeable informant in significantly more switch trials than children who did not meet criterion, we looked more closely at their performance to explore potential age effects on the use of reasoning-based SLSs. If children were using explicit reasoning about the informants' knowledge, we would expect to see more selections of the knowledgeable informant. Alternatively, if children were relying on simple associations, or inferred rules, that were successful in the pre-switch trials but not based on an understanding of the informants' potential to provide information, we would expect to see informant selections that were closer to chance. To address whether age affected the likelihood of using explicit reasoning to determine who was the appropriate informant, we constructed a GLMM for successful selection of the knowledgeable informant in each switch trial for participants who met criterion. The GLMM had fixed effects of age, the presence of frames, the interaction between these variables, and a random intercept of participant ID. The results showed a significant main effect of age ($b = 1.63$, $SE = 0.51$, $z = 3.18$, $p = .001$), which suggests that, as expected, older children selected the knowledgeable informant significantly more often than younger children. This pattern of results could indicate that some of the younger children who had succeeded in the pre-switch trials had done so by forming a simple association, or inferring a rule, but that older children were more likely to have explicitly reasoned the informants' knowledge based on their perceptual access. There was no main effect of the presence of frames ($b = -0.05$, $SE = 0.20$, $z = -0.27$, $p = .786$) or interaction between age and the presence of frames ($b = 0.16$, $SE = 0.42$, $z = 0.38$, $p = .706$).

Alternative Approaches to the Task

Of the 52 children who met criterion, 65% did so after only five pre-switch trials, while the rest required between six and ten trials to meet criterion and move on to the switch phase (see Table 1). Meeting criterion after only five pre-switch trials *and* continued success in the switch trials could indicate reasoning of the informants' knowledge from the outset of the task. We compared children who met criterion in only five pre-switch trials with children who required more than five pre-switch trials to assess whether the number of trials required to reach criterion influenced children's performance in the switch trials. A GLMM was built for successful selection of the knowledgeable informant in each switch trial with fixed effects of age, whether criterion was met in five or more than five pre-switch trials, the presence of frames, the interactions between these variables, and a random intercept of participant ID. The GLMM revealed a significant main effect of age ($b = 1.45$, $SE = 0.50$, $z = 2.88$, $p = .004$), indicating that older children selected the knowledgeable informant more often than younger children. A significant main effect of the number of pre-switch trials required to meet criterion ($b = 0.80$, $SE = 0.26$, $z = 3.10$, $p = .002$) suggested that children who met criterion after only five pre-switch trials selected the knowledgeable informant more often than children who required more than five pre-switch trials to reach criterion. Again, there was no evidence of an effect of the presence of frames (b = −0.11, SE = 0.24, z = −0.46, p = .648) or of any interactions between the fixed effects ($p \geq .074$).

One-sample *t*-tests were used to compare the within-participant switch trial performance of each group to chance (50%). The results showed that children who met criterion in five trials selected the knowledgeable informant in an average of 4.59 switch trials (see Table 2), this was significantly above chance $t(33) = 20.0$, $p < .001$, one-tailed. Children who required more than five trials to reach criterion selected the knowledgeable informant in an average of 3.56 switch trials, also significantly above chance $t(17) = 3.24$, $p = .002$, one-tailed. An independent-samples *t*-test showed that the difference in the number of knowledgeable informant selections between these two groups was significant $t(50) = 3.75$, $p < .001$, one-tailed.

To investigate whether children who met criterion after five or more than five trials performed differently to the children who did not reach criterion, we compared the within-participant performance of each group. A one-way ANOVA revealed a significant difference between the switch trial performance of children who did not reach criterion, children who met criterion in only five pre-switch trials, and children who met criterion after more than five pre-switch trials $F(2, 106) = 24.15$, $p < .001$ (see Figure 4). Post hoc Tukey HSD tests showed that children who met criterion after five pre-switch trials (M = 4.59 trials, SD = 0.61, N = 34) selected the knowledgeable informant significantly more often than both children who met criterion after more than five pre-switch trials (M = 3.56 trials, SD = 1.38, N = 18, p = .006) and children who did not reach criterion (M = 2.89 trials, SD = 1.26, N = 57; $p < .001$). However, there was not a significant difference in performance between children who met criterion after more than five pre-switch trials and children who did not reach criterion (p = .08).

Figure 4

Number of Switch Trials in Which Children Selected the Knowledgeable Informant by the Number of Pre-Switch Trials Required to Meet Criterion

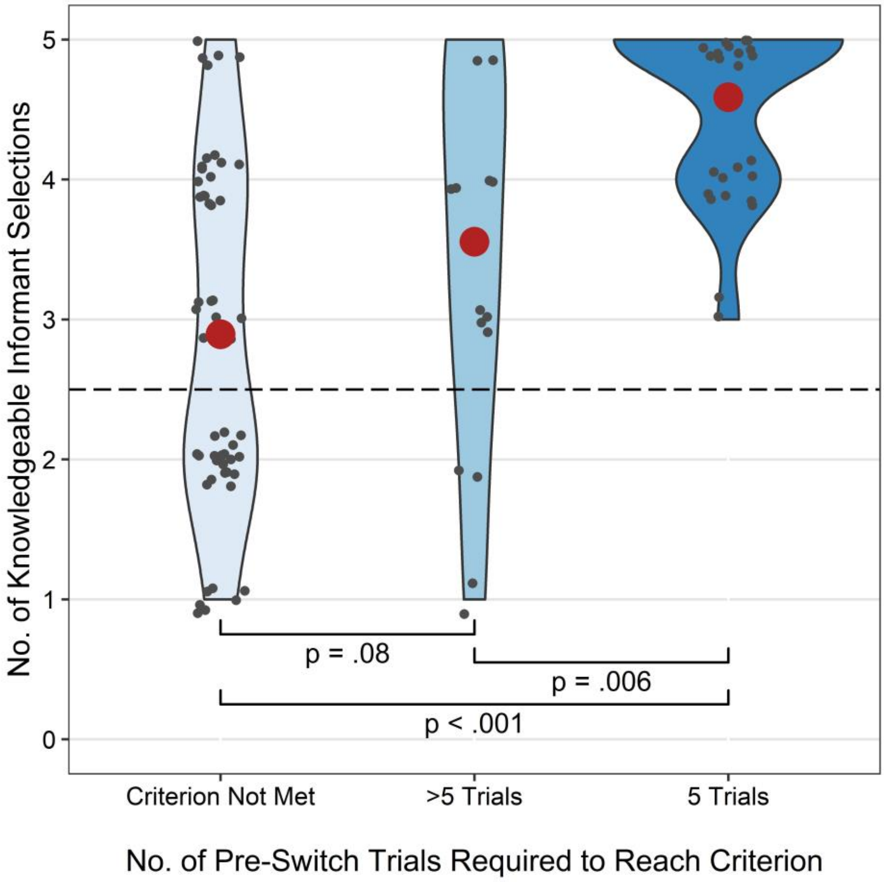

Note. Children who met criterion after only 5 trials performed significantly better than children who took more than 5 trials to reach criterion and children who did not reach criterion. Small (black) points show individual participants' performance, large (red) points indicate group means. Brackets indicate differences between groups. Dashed line indicates chance performance.

Explicit Verbal Reasoning

We examined children's responses to the explicit reasoning question regarding how they were choosing who they wanted to show them where the Lego was hidden. Children were asked this question after they had completed the switch phase, so each child provided a single response. Responses to the reasoning question were categorised into four levels: 0) no response; 1) non-reasoned responses, i.e., responses that did not relate to the task (e.g., *"because I used telepathy, cause I watched a video about it I just used my brain"),* or comprised of single words or gestures; 2) reasoned but incorrect responses, i.e., explanations that showed evidence of explicit reasoning but the motivations were incorrect (not related to visual perceptual access) or the answer was not sufficient to determine full correct reasoning (e.g., *"I found a pattern, it switched from Spots to Stripes from Spots to Stripes"*), and 3) reasoned correct responses, i.e., the explanation provided clear reasoned evidence of explicit task understanding with reference to visual perceptual access (e.g., *"The person that had the blue frame could see what you were doing, the person who was turned around you couldn't see, the person with the yellow frame couldn't see"*). When asked to justify how they made their choices 41% of children provided reasoned correct responses.

We investigated whether children's verbal reasoning was predicted by age, meeting the proficiency criterion, or switch trial performance (number of switch trials on which each participant selected the knowledgeable informant) by conducting ordinal regressions (using the ordinal package in R, Christensen, 2019). Firstly, children's reasoning responses were submitted to an ordinal regression with fixed effects of age, meeting criterion and the interaction between these variables. This model indicated significant main effects of both age (b = 1.98, SE = 0.44, z = 4.50, p < .001) and meeting criterion (b = 0.83, SE = 0.21, z = 3.97, p < .001; see Figure 5). These results suggest that older children and children who met criterion provided better reasoned responses than younger children and children who did not meet criterion, respectively. There was no evidence of a significant interaction between age and meeting criterion (b = 0.83, SE = 0.42, z = 1.03, p = .305).

Figure 5

Proportion of Participants Providing Each Category of Explicit Verbal Reasoning Responses by Age in Years and Whether Children Met Criterion or Not

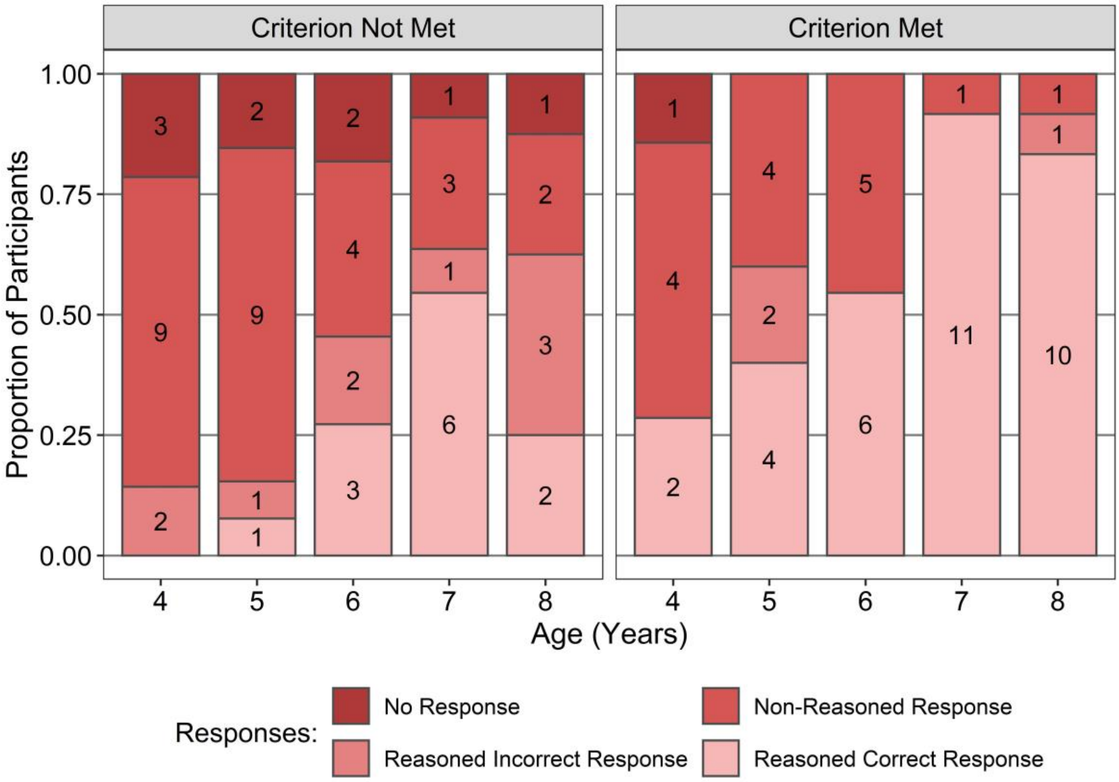

Note. Values presented on bars refer to the number of participants each proportion represents.

To explore the relationship between switch trial performance and verbal reasoning we conducted a further ordinal regression. This model included fixed effects of age, met criterion, switch trial performance and the interactions between these variables. The results revealed a main effect of switch trial performance (b = 1.30, SE = 0.26, z = 5.00, p < .001) indicating that children who selected the knowledgeable informant on a higher number of switch trials also provided better reasoned responses to the explicit reasoning question (see Figure 6). In this model there were no significant main effects of age (b = −1.93, SE = 1.81, z = −1.065, p = .287) or meeting criterion (b = −1.42, SE = 0.95, z = −1.50, p = .133). However, the model revealed a significant interaction between switch trial performance and age (b = 1.01, SE = 0.47, z = 2.14, p =

.032). This interaction is likely to be, at least in part, the product of older children's switch trial performance being higher than that of younger children. However, the effect of switch trial performance (i.e., children who performed well in the switch trials providing better reasoned responses) appears to be enhanced by age. Given the absence of a main effect of age, this interaction suggests that older children who performed well in the switch trials were more likely to provide a better reasoned response than younger children who also performed well. This pattern of results suggests that the successful switch trial performance of both older and younger children was driven by reasoned understanding of the informants' perceptual access, though this effect was greater in older than younger children. Thus, the successful identification of the appropriate informant appears to have been driven by explicitly metacognitive SLSs.

Figure 6

Proportion of Participants Providing Each Category of Explicit Verbal Reasoning Responses by Number of Knowledgeable Informant Selections

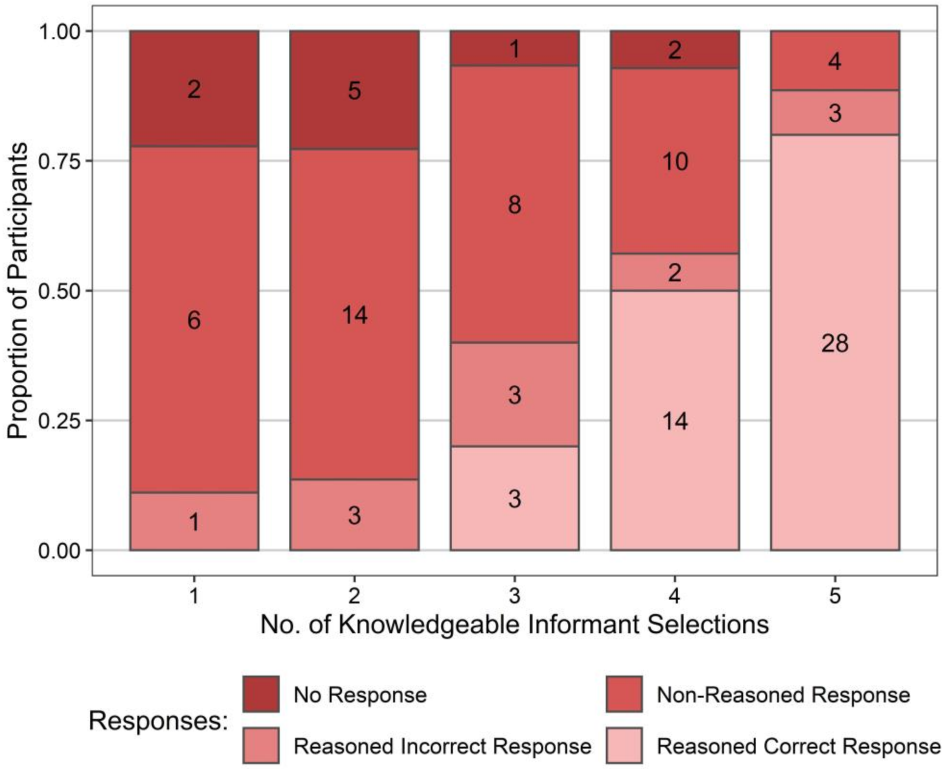

Note. Values presented on bars refer to the number of participants each proportion represents.

Reasons for Not Reaching the Proficiency Criterion

We also explored the possibility that the 52% of children who did not meet criterion were employing alternative yet consistent strategies regarding informant selections. There were a number of potential strategies that could have interfered with children's formation of an association, or inference of a rule, in the number of trials available to them. For example, children may have preferred a particular informant, they may have alternated their choices between the sides, or they may not have used any strategy at all. A chi-square test of independence showed no evidence of an overall preference for either informant (χ^2 (1) = 0.18, p = .675) or of any particular age group (years) having a preference for either informant (χ^2 (4) = 2.10, p = .717). To explore whether children had any preferences for copying or not copying the side they selected on the previous trial we conducted further chi-square tests of independence. These revealed no evidence of preference for either copying or not copying the side selected on the previous trial (χ^2 (1) = 0.02, p = .895) and no evidence of that any age group preferred one of these choices (χ^2 (4) = 3.15, p = .533). These results suggest that it is unlikely that children who did not reach criterion were using a consistent alternative strategy for informant selections.

Discussion

To make best use of learning from others, learning should be targeted towards those who possess the most relevant knowledge or experience. Adult humans' skill at identifying appropriate sources of social information and tracking 'who knows' has been attributed to the use of explicitly metacognitive SLSs (Dunstone & Caldwell, 2018; Heyes, 2016c). Such capacities offer a great deal of potential in terms of revolutionising the strategic use of social information, relative to the use of implicit SLSs. Despite this, there have been few attempts to explore the presence or emergence of these abilities in children. It has been argued that explicitly metacognitive SLSs are what set human learning apart from that of animals and that such a capacity may be responsible for distinctively human cumulative culture. Thus, the current study set out to track the development of the capacity to use explicitly metacognitive SLSs in 4- to 8-year-old children relative to an alternative learning process (their reliance on implicit SLSs).

To differentiate the SLSs that children have the capacity to use, we investigated whether there were any age-related changes that could reflect a switch from the use of implicit SLSs to explicitly metacognitive SLSs. To do this we employed a paradigm more commonly used in non-human primate research. We exposed children to iterated trials in which they could reach a proficiency criterion either by learning an association, inferring a rule, or making a causal interpretation based on mentalistic reasoning about others' knowledge. Following a switch in the scenario, to continue to be successful, children were required to explicitly reason others' perceptual access to a critical event and assess their suitability as a source of information. The switch was intended to generate a conflict between the response favoured by simple associations and the response favoured by a mentalistic understanding of why another's behaviour is informative. If children were reasoning others' suitability as a source of knowledge based on their perceptual access, we expected that they would continue to be successful on switch trials. However, if children were relying on implicit SLSs, we expected that they would struggle to continue to be successful in switch trials, as use of associations, or inferred rules, that were predictive of success in the pre-switch trials would no longer be appropriate. Thus, high performance in the switch trials would be consistent with use of reasoning-based explicitly metacognitive SLSs.

The results of this study demonstrate that it is possible to differentiate the use of implicit and explicitly metacognitive SLSs empirically. Overall, we found that the children who reached the proficiency criterion selected the knowledgeable informant in significantly more switch trials than children who did not reach criterion. This suggested that at least some of the children who met criterion were explicitly reasoning about the informants' perceptual access to identify which informant was knowledgeable. Children who did not meet criterion selected the knowledgeable informant significantly less often than children who met criterion, suggesting that they were less likely to have reasoned the informants' perceptual access, or recognised the value of the information that they could provide. The behaviour of these children was not wholly consistent with either having formed an association, inferred a rule, or used an appropriate reasoning-based strategy, any of which would have been expected to generate success in the pre-switch phase. Accordingly, their informant selections in the switch trials were also less selective. To investigate whether there were age-related changes that indicated differences in the learning processes underlying children's selections, we analysed the switch trial performance of those who met criterion. Our discussion will focus on the switch trial performance of children who reached the proficiency criterion, though first, we will consider explanations for why so many children failed to reach the proficiency criterion.

It is not trivial that 52% of children did not reach the proficiency criterion. It is important to note, however, that we were not expecting all children to do so. The task required children to form an association, or infer a rule, very quickly. In order to get five consecutive correct trials, a successful rule would need to be inferred, or an association firmly established, by only the sixth trial, given that we had set a cut-off point at ten trials. This was quite a high expectation for young children. However, it would not have been practical to have children complete more pre-switch trials. Children can only stay focused on a single task for a limited time, and it was critical for our results that they retained motivation to complete the final switch phase of the task. It is of course possible that some of the children who did not reach criterion had indeed formed an association, or inferred a successful rule, by the later pre-switch trials beyond the point at which they could have reached the proficiency criterion. We found no evidence that the children who did not reach criterion were using any alternative yet consistent strategies in the pre-switch trials upon which to base their informant selections. Therefore, the failure to form an association, or infer a rule, within the available number of trials was likely due to random selections. Similarly, there is no evidence that these children had a causal understanding based on mentalistic reasoning about the informants' knowledge. It is possible that the no-risk nature of the pre-switch trials and that children were rewarded after every selection may have resulted in a low motivation to identify the knowledgeable informant. Always receiving positive feedback in the pre-switch phase may have limited children's motivation to identify the knowledgeable informant. Yet, it was arguably beneficial that the knowledgeable informant was the only person to provide feedback regarding the location of the reward. It reduced the potential for problems to arise as a result of the ignorant informant either 'guessing' correctly or being consistently incorrect. Children may have become confused in either such case, as guessing correctly may have confounded the formation of associations and being wrong 100% of the time is very unlikely. If we had chosen for the ignorant informant to provide feedback in the pre-switch phase, we would have required a much larger number of pre-switch trials for children to learn an association or infer a rule. In contrast, in the switch phase we chose to have the ignorant informant always reveal the non-target location if they were selected so that there was a clear consequence for choosing the ignorant informant.

Age-Related Transition From Implicit to Explicitly Metacognitive SLSs

We found that of those who met criterion, older children sought information from a knowledgeable informant over an ignorant informant significantly more often than younger children in the switch trials. This developmental trend was consistent with our prediction that there would be age-related differences in the SLSs underpinning success in the pre-switch trials. We propose that these results reveal a relatively late age-related transition from the use of implicit SLSs to the use of explicitly metacognitive SLSs. Such a finding is also consistent with evidence from a previous study (Blakey et al., 2020, Chapter 3) that showed that children's information seeking was not driven by reasoned understanding of the information required to identify appropriate sources of information until 7 to 8 years of age. That older children in the present study selected the knowledgeable informant more often suggests that they had the capacity to assess others' suitability as sources of knowledge based upon their perceptual access (as inferred from personal experience). In contrast, at least some of the younger children who had succeeded in the pre-switch trials appeared to have done so using visually salient, yet superficial, cues such as colour or side to learn a simple association or infer a rule. However, such associations or rules were no longer appropriate in the switch trials. This offers an explanation for why younger children made fewer selections of the knowledgeable informant in the switch trials despite reaching criterion. Reasoning-based information seeking therefore appears to be cognitively challenging, and its relatively late development indicates that the metacognitive capacities required may preclude younger children from identifying the most appropriate source of social information. Therefore, if younger children appear to be relying upon imperfect implicit SLSs in place of explicitly reasoning others' perceptual access, it may be that they have not yet developed the necessary metacognitive capacities to support explicitly metacognitive SLSs. Specifically, younger children might not yet understand that perceptual access grants privileged knowledge, or if they do, they may not appreciate that this renders an individual a valuable source of information. Though we propose that adults use explicitly metacognitive reasoning, we recognise that once adults have had extensive experience, they may switch to using efficient implicit rules (Saling & Phillips, 2007) even if they can, in theory, explicitly understand. However, we do not believe that is the case for the children in this study due to the very limited number of trials to which children were exposed.

Children who met criterion after only five pre-switch trials selected the knowledgeable informant significantly more often in the switch trials than both children who took more than five pre-switch trials to reach criterion, and children who did not reach criterion. However, there was no difference between children who did not reach criterion and children who met criterion after more than five trials. These results suggest that the two groups of criterion-passers were employing different SLSs. Those who met criterion after only five trials continued to select the knowledgeable informant in the switch trials. As outlined above, we believe this is indicative of reasoning the informants' knowledge on the basis of their perceptual access. In contrast, those who required more than five trials to reach criterion were relatively less successful in identifying the knowledgeable informant in the switch trials, which could suggest a reliance on implicit SLSs. We do not claim that all of the children in this group were relying on implicit strategies, as it is possible that some were reasoning the perceptual access of the informants (albeit dependent on task experience to make this inference) while others may have reached criterion by forming an association between superficial visual cues and success. Such disparity could explain why we see an average performance that is greater (though not significantly so) than children who did not reach criterion but lower than those who met criterion in the minimum number of trials.

In the final two switch trials we removed the necessity to understand the properties of the frames by removing them and affording perceptual access via the knowledgeable informant facing towards, and the ignorant informant facing away from, the critical event. Removing the frames also allowed exploration into whether participants were truly responding to the informants' perceptual access, or whether they had simply made the connection between success and the frame, or frame colour, that they could have been continuing to use as a cue to predict success. We found no effect of the presence of the frames in any of our analyses suggesting that children who were reasoning the perceptual access of the informants continued to do so when the frames were removed. This also suggests that these children were unlikely to have been making selections based on an association, or rule, related to the colour of the frame. Equally, children who were not reasoning perceptual access by way of the frames also appear not to have been reasoning the perceptual access of the informants when perceptual access was portrayed in the more visually salient manner of the informants facing and not facing the critical hiding event. If we had found that children performed better when perceptual access was afforded by the informants facing or not facing the critical event, then it may have been that children recognised the value of the information offered by others on the basis of perceptual access but were unable to use their personal experience of the properties of the frames to determine others' perceptual access and thus knowledge state. Yet this is not supported by the results; rather, it appears that those children who recognised the value of perceptual access as a source of knowledge were able to use their personal experience of the frames to identify which of the two informants had the desired knowledge, showing a similar capacity to that observed in human adults (Teufel et al., 2009).

Due to the nature of the methods we used it was not possible to determine the reasons for lower success of either the children who did not reach criterion, or those who struggled to switch from using implicit SLSs to using explicitly metacognitive ones. It is possible that they lacked the necessary conceptual understanding regarding the perceptual access of the informants. Alternatively, they may have possessed this understanding but could not use this to deduce which of the informants could provide valuable information. As highlighted in the introduction, it may be that younger children are able to identify others' perceptual access but may not yet recognise the value of the information that that access grants. Studies such as Robinson et al. (2011) have found that when required to choose one of two puppets (knowledgeable vs ignorant) young children (4;3- to 5;2-year-olds) were no better than chance at asking the puppet with the relevant experience. This supports the theory that younger children do not recognise the potential value to themselves of others' knowledge, despite finding it relatively easy to report others' knowledge based on perceptual experience. Due to the apparent cognitive challenge associated with actually using such information, the requirement to do so in this task may have precluded many children from selecting the appropriate informant. Therefore, it is possible that the children in the current study may have been capable of reporting the informants' knowledge state had we asked them to do so. We believe that our findings highlight just how challenging using explicit reasoning about others' knowledge in the context of social learning is likely to be. Assumptions about the use of such strategies in the social learning literature should therefore be made and interpreted with extreme caution, especially with regards to the capacities of young children or animals.

The explanation above suggests that it is possible that the younger children do understand who has perceptual access, but simply fail to see the relevance of this as a means to address their own ignorance. However, it also remains possible that the younger children's failure to show the necessary flexibility of response during the switch trials could have arisen for reasons completely unconnected to conceptual understanding. In particular, this could potentially be explained as a result of limited capacities for inhibition, leading to perseverative errors which would cause children to persist with a previously successful response in spite of knowledge of an updated rule (e.g., see examples from the Dimension Change Card Sort Test, Diamond, 2013; Zelazo et al., 1996). However, we would argue that children who recognised the significance of perceptual access should in fact give *this* rule primacy, over alternative interpretations focussed on one or other (or some combination) of the other cues which were only arbitrarily linked to success. Thus, if children were perseverating with a response other than one driven by the perceptual-access rule, this seems to suggest that at the very least they were attending to cues other than those associated to perceptual access.

Explicit Verbal Reasoning

Our analysis of children's responses to the explicit verbal reasoning question revealed that older children and children who met criterion were rated as having given better reasoned responses. However, when switch trial performance was included as a predictor, the effects of both age and criterion were eclipsed, and we discovered that children who selected the knowledgeable informant more often in the switch trials were also more likely to have been rated as having provided appropriate verbal reasoning for their choices. Successful identification of the appropriate informant appears to have been driven by explicitly metacognitive SLSs regardless of age or whether children met the proficiency criterion. Children's responses to our explicit reasoning question offered an insight into their understanding of the benefits of perceptual access on others' knowledge. The results gained from children's responses are consistent with the interpretation that some children were indeed employing reasoning-based understanding regarding others' perceptual access to assess their suitability as a source of social information. This was true irrespective of age, supporting the explanation that selections of the knowledgeable informant were driven by an understanding of why another is an appropriate source of knowledge. Similarly, Butler et al. (2020) recently reported that 6- and 7-year-old children's ability to select information on the basis of claims that have been verified (via perceptual access) is related to a developing ability to offer appropriate justification for why they made those selections.

Conclusions

The results of this study indicate that it is possible to differentiate implicit and explicitly metacognitive SLSs empirically. The methods we employed, drawn from existing paradigms in the non-human primate literature, facilitated investigation into the developmental trajectory of the transition from use of one process to the other. We found that, of all those who met the proficiency criterion, older children were more likely to seek information from the knowledgeable informant in the switch trials, indicating that they may have taken a different approach to the task. The pattern of results suggested that some of the younger children who had succeeded in the pre-switch trials had done so by forming a simple association, or inferring a rule, based on superficial visual cues. Older children were more likely to have explicitly reasoned others' perceptual access. They appeared to have used personal experience to infer others' perceptual access and thus assess their suitability as sources of knowledge. This ability was rare in younger children. Rather, younger children's selections of both the knowledgeable and ignorant informants suggested that they were continuing to rely upon imperfect implicit strategies learned during the pre-switch trials that were no longer appropriate. Moreover, we discovered that children who sought information from the knowledgeable informant more often also provided more appropriate verbal reasoning for their choices, irrespective of age. This supports the interpretation that informant selections were driven by an understanding of why another is an appropriate source of knowledge. This disparity in children's approaches to selecting informants is consistent with previous findings (Blakey et al., 2020, Chapter 3) that indicated an age-related transition from the use of implicit heuristic biases to explicitly metacognitive

SLSs. The similar developmental trajectory identified in the current study suggests that the capacity to reason about others' knowledge develops relatively late in childhood. This late development supports Heyes's (2016c) proposal that such capacities are experience-dependent. The developmental trend captured within our results may also be indicative of a significant cognitive challenge associated with employing explicitly metacognitive SLSs. Hence, these metacognitive capacities may preclude younger children from recognising the value of others' knowledge. Therefore, we argue that being able to fully benefit from others' knowledge appears to be dependent, at least in part, on the ability to assess the others' suitability as sources of knowledge based upon their perceptual access to desired information. The relatively late development of explicitly metacognitive SLSs suggests that such capacities are unlikely to be observed in animals, offering credence to proposals that such capacities may be involved in distinctively human cumulative culture. If human cumulative culture is indeed underpinned by explicitly metacognitive SLSs, then the degree of flexibility afforded by an explicit understanding of others' mental states, with regards to assessing others' suitability as informants, may offer the significant advantage in social information use that drives human cultural evolution in a way not seen in non-human species.

Chapter 5:

General Discussion

The overall aim of this thesis was to investigate cognitive mechanisms that drive the development of distinctively human cumulative culture. In particular, the studies reported in the preceding three chapters were concerned with investigating fundamental questions related to the development of cognitive capacities for social information use. Their premise was to look for age-related changes in children's social information use in response to particular cognitive challenges that had been identified as being potentially relevant in the context of real world cases of cumulative culture.

While recent evidence has highlighted some capacity for cumulative culture in non-humans (Jesmer et al., 2018; Sasaki & Biro, 2017; Schofield et al., 2018) its expression appears to be context-specific and relatively restricted compared to human adults (Caldwell & Millen, 2008b; Mesoudi & Thornton, 2018; Tomasello et al., 1993). A variety of sociocognitive mechanisms have been proposed as necessary prerequisites (Boyd & Richerson, 1996; Dean et al., 2014; Lewis & Laland, 2012; Tennie et al., 2009). However, many of the proposed prerequisite mechanisms, such as imitation and selective social learning, are common to both humans and non-humans. Therefore, they cannot fully account for the differences between these populations with regard to the capacity for cumulative culture. This thesis adopted a developmental approach, examining the ontogeny of sociocognitive mechanisms, to assess whether there is a particular reason, from a cognitive perspective, for this discontinuity between humans and non-humans.

In this thesis I have investigated the feasibility of an alternative proposal that attributes the distinctiveness of human cumulative culture to the capacity to employ explicitly metacognitive SLSs (Dunstone & Caldwell, 2018; Heyes, 2016c). This proposal suggests that these reasoning-based strategies could enable learners to recognise the potential value, to themselves, of information provided by others. The capacity to recognise the relevance or value of social information would likely offer a significant advantage in effective social information use. Given that human adults are assumed to have the capacity for explicitly metacognitive SLSs, and that the social information use of both non-humans and young children is proposed to be driven by implicit SLSs (Heyes, 2017), it follows that the capacity for explicitly metacognitive SLSs develops during the course of childhood. Therefore, to investigate the developmental trajectory of the capacity for explicitly metacognitive SLSs this thesis examined age-related changes in children’s social information use. Changes in children’s use of social information could indicate that they are employing different cognitive capacities at different ages. Therefore, I designed the studies in such a way that it would be possible to establish whether children’s social information use was driven by reasoned understanding of its value (explicitly metacognitive SLSs) or whether it could be explained by more implicit adaptive heuristics (implicit SLSs).

Looking at whether age-related changes in social information use coincide with advances in cognitive development could offer an insight into the underlying mechanisms that are required and help to explain why cumulative culture appears to be restricted to humans. That is, if we find particular ages at which children make significant advances in their appropriate social information use and these occur at a similar age to the development of particular cognitive capacities, then those capacities may be necessary prerequisites. If these capacities emerge late in development, this may indicate that they are beyond the capabilities of non-humans and therefore could offer an explanation for the distinctiveness of human cumulative culture. Indeed, I proposed that the capabilities of younger children in terms of social information use are likely to also be shared with non-humans.

In this chapter I will present an overview of the key findings from each of the studies presented. I will then discuss these findings more broadly, relating them to the existing literature and highlighting the novel contribution they make. Finally, I will consider some of the limitations that could be associated with the approaches taken in this thesis and outline avenues for further research.

Taking Account of Others’ Goals

The first cognitive challenge that I identified as being potentially relevant for social information use was the ability to account for others’ goals. Having the ability to understand and take into account others', potentially conflicting, motivations is likely necessary for recognising whether the information they have provided is valuable and relevant for achieving one's own goal. Human adults regularly use their understanding of others' goals, relative to their own, to identify and use the most relevant social information (Caldwell, 2018; Vélez & Gweon, 2019), yet there is little evidence of such capacities in children or non-humans. Thus, investigating the developmental trajectory of these abilities and the cognitive challenges they pose to young children (and by logical extension non-humans) could advance our understanding of the cognitive capacities that are necessary for human-like cumulative culture.

Therefore, in chapter two I investigated age-related changes in 3- to 7-year-old children's ability to use social information by taking into account others' goals relative to their own. In particular, this study considered whether children's abilities to interpret and use social information were influenced by the degree to which others' goals aligned with their own. To do this I employed a novel paradigm which assessed children's understanding of others' goals relative to their own, their ability to use this knowledge to interpret the outcomes of social demonstrations, and their ability to apply this understanding in order to use the social information effectively. Crucially, in contrast to previous social learning paradigms, this study specifically compared cases in which the goals of the demonstrator and the participant were aligned with cases in which their goals were not aligned. In this task children observed a demonstrator select a capsule from one of two buckets, the demonstrator peeked inside and chose to accept or reject the capsule. Children could then choose to copy the demonstrator's selection or shift and choose a capsule from the alternative location. To interpret the outcome of the demonstration children had to take into account the demonstrator's goal and recognise that the demonstrator had made an informed choice (based on access to information). To make best use of the available information, children needed to recognise the value of the information that had been provided in the demonstration in respect to their own goal while also taking into account the goal of the demonstrator. Therefore, even if children understood the task goals and correctly inferred the outcome of the demonstration, this would not necessarily mean that they would be able to use this information appropriately.

The results revealed that children were capable of understanding and representing the demonstrators' goals from around 4 years old and could reliably interpret the outcome of the demonstrators' behaviour from around 6 years old. This was true whether the demonstrator's goal conflicted with their own or not. Therefore, contrary to expectations, the results indicated that the degree to which others' goals aligned with their own did not influence children's ability to understand the goals, or their ability to interpret the outcome of others' behaviour. Finding what appeared to be an age-related jump from understanding others' goals to reliably *interpreting* the demonstration outcome, indicated that this ability was dependent on advancing cognitive development. The age-related change did indeed correspond with advances in cognitive development reported in the literature. Previous research claims that children do not demonstrate understanding of how knowledge is formed with consideration of the role of access to information until around 6 years (Kloo & Rohwer, 2012). Therefore, 6- and 7-year-olds' more reliable interpretations of the demonstration outcome may have been related to a cognitive advancement with regard to understanding that the demonstrator had had perceptual access to the contents of the capsule and was therefore informed.

The results also showed evidence of the expected age-related improvement in children's appropriate social information use. Here too, the alignment of the demonstrators' and participants' goals was not found to influence performance. However, even the oldest children's performance was not particularly high, suggesting that the task paradigm was more challenging than anticipated. At first glance the age-related improvement in appropriate information use could be seen as evidence for explicitly metacognitive SLSs. However, the improvements were restricted to cases in which the appropriate response was to shift. Thus, this pattern of results suggested that even the older children may not have been using explicitly metacognitive SLSs. By contrast, the near ceiling performance of a small group of adult participants provided a useful comparison for what to expect from a participant who is expected to have the capacity for explicitly metacognitive SLSs. Thus, the distinct difference between the performance of the oldest children and the adults supports the interpretation that some children were not using explicitly metacognitive strategies. In the absence of the capacity for explicitly metacognitive SLSs, children's choices may have been driven by implicit biases or motivations towards exploration. Indeed, an alternative explanation for this pattern of responses suggests that they might be an artefact of the binary choice task. Recent evidence indicates that children often choose to explore even in situations in which they might benefit from exploitation (Blanco & Sloutsky, 2020). This may be due to a trade-off between explore and exploit motivations (Atkinson et al., 2020), which overall, in the context of this task, would favour shift responses. If so, it is possible that children may have recognised the value of the social demonstration, but the motivation to explore may have overridden the motivation to achieve the goal, making it difficult to ascertain the reason for particular decisions.

The findings of this study make some key contributions to the literature. First, for children aged 4 years and older the degree to which others' goals align with their own does not appear to influence children's ability to understand others' goals, their ability to interpret the outcome of others' behaviour, or their ability to use social information. Second, it appears that taking others' goals in account at all in appropriate social information use is cognitively challenging. Even at 7 years old children's performance did not demonstrate an adult-like ability to account for others' goals in information use. Thus, even if children have the ability to interpret the content of social information, they may still struggle to recognise its value to themselves.

Overall, chapter two provided some evidence to support the hypothesis that explicitly metacognitive SLSs are cognitively challenging and develop relatively late in childhood. The significant challenge associated with the ability to use social information about others' goals also provides some support for the hypothesis that it depends on cognitive developments that could be beyond the capabilities of non-humans.

Appropriate Social Information Seeking

While children's use of social information can indicate to some degree their understanding of it, it is difficult to determine from social information use alone whether children have recognised its value and are indeed employing explicitly metacognitive SLSs. Looking at how they seek out or select information to attend to offers more potential in terms of establishing whether they recognise its value. Indeed, in reality, learners are rarely passive recipients of information that is relevant for achieving a specific goal. Rather, learners usually have to actively seek out information from a variety of potential others who are in a position to provide useful information. Therefore, chapters three and four examined the development of children's ability to seek out information from appropriate social sources.

Recognising What Information is Required

In chapter three I investigated another cognitive challenge, this time related to the ability to recognise what information is required to achieve a particular goal. Being able to recognise what information is needed to solve a particular problem may be key to appropriate social information seeking as it might enable learners to identify who could provide relevant and useful information. Furthermore, having an understanding of what information is needed also suggests that learners would know what to do with the information once they have acquired it (Baldwin & Moses, 1996). Investigating children's social information seeking offers the potential to determine the age at which children develop the capacity for explicitly metacognitive SLSs. As well as this, it might provide an insight into the alternative learning processes (i.e., implicit SLSs, such as model-based biases) children rely on prior to its development.

The study presented in chapter three looked at age-related changes in 3-to-8-year-old's ability to seek out and use social information appropriately, based on an understanding of what information was needed to solve a particular problem and who could provide that information. To do this I looked at who children chose to seek information from, and how they then used that information, as a means to determine whether their selections were based on an understanding of the potential value of the information the model could provide (in this case, access to relevant information), or on superficial, yet salient, demonstrator characteristics such as age or gender. The task was designed such that relying on model-based biases (implicit SLSs) would be inappropriate. So, if children were seeking out information appropriately, based on a reasoned-understanding of the value of the information that others could provide, then model-based biases should be overridden in favour of reasoning-based strategies (explicitly metacognitive SLSs). To assess this, in each trial children were presented with a novel problem in which they had to choose between two coloured keys to unlock a box locked with a coloured padlock. Before selecting a key, children had to select one of four potential demonstrators (an adult and a child of each gender) to watch a video of them unlocking a box. One target demonstrator faced the same problem (combination of padlock and keys) as the child, while three non-target demonstrators faced equivalent but different problems. After watching a successful or unsuccessful demonstration children selected one of the keys to try and unlock their box. Appropriate information seeking required children to reason about the information they needed to solve the problem and which of the demonstrators could provide that information (i.e., who had the same problem). Manipulating the age and gender of the target and non-target demonstrators across trials enabled exploration of model-based biases in children who were not seeking social information appropriately.

As expected, the results revealed an age-related improvement in children's appropriate information seeking and appropriate information use. As predicted, older and younger children appeared to approach the task differently. Older children (7 and 8 years) appeared to have been able to reason about the information that they needed, and to have recognised that the target demonstrator faced the same problem as themselves, so the information that they could provide was relevant. That older children also knew what to do with the information that they got, as reflected in their high rate of appropriate information use, suggested that they may have understood the potential value of the information prior to selecting a demonstrator. By contrast, younger children (3 to 6 years) appeared to rely on imperfect model-based biases for gender congruence when they struggled with information seeking. These children tended to select target and non-target demonstrators who were the same gender as themselves. So too, their less effective use of social information suggested that they had not recognised what information was needed, had not used reasoning-based strategies to select the demonstrator, and therefore did not know what to do with the information that they received.

Interestingly, despite relatively low rates of appropriate information seeking, the appropriate information use of 6-year-olds who selected the target demonstrator was comparable to that of 7- and 8-year-olds. This suggested that those 6-year-olds who selected the target demonstrator might have been driven by reasoned understanding of the relevance of the information. The developmental trajectory of appropriate information seeking, and use, pointed to cognitive advances in explicit metacognition being involved in the ability to recognise the value of social information. Such advances are related to the ability to reflect on one's own knowledge and, importantly, the ability to recognise gaps in one's knowledge to identify learning needs. This ability has been found to emerge at around 6 years old (Rohwer et al., 2012).

Overall, the results of chapter three indicated that children's capacity to use explicitly metacognitive SLSs develops relatively late in childhood. The developmental trajectory suggested that, in this context, the transition from using implicit SLSs to using explicitly metacognitive SLSs occurred between 6 and 7 years old. The findings also suggest that appropriate information seeking might have been overlooked as a significant challenge involved in fully benefitting from others' knowledge. Importantly, in contrast to the study presented in chapter two, the methods employed in this study made it possible to determine with greater certainty whether children were more likely to have reasoned the value of social information, or whether they had instead made selections based on superficial yet salient characteristics such as age or gender.

Identifying Who Knows

Chapter four aimed to address whether children can use explicitly metacognitive SLSs in cases in which it is necessary to consider others' mental states. In particular, chapter four looked at the developmental trajectory of 4- to 8-year-old's ability to identify potential informants' suitability as sources of information based on an understanding of the informant's knowledge, which in this case was privileged perceptual access.

This study used a novel methodology adapted from studies of social understanding in non-human primates. Children were first given experience of two different frames, a semi-opaque one which permitted perceptual access and an opaque one which did not. From a distance the opacity of the frames looked the same – though they were distinguishable by the colour of the frame. They were then exposed to repeated trials in which they had to choose which of two informants to ask for information about a critical event during which the knowledgeable informant was behind the semi-opaque frame and the ignorant informant was behind the opaque frame. The purpose of these trials was to give children the opportunity to learn by forming an association or inferring a rule based on a superficial, yet salient, visual cue such as the side of the table. Following these trials there was a 'switch' in the task scenario. Switching generated a conflict between the response favoured by simple associations and the response favoured by a mentalistic understanding of why another's behaviour is informative (i.e., who had perceptual access and therefore who was knowledgeable). The switch occurred once a proficiency criterion (five consecutive correct responses) was reached, or after 10 trials. Following the switch, continued use of an association, or inferred rule, would result in imperfect responses, leading to lower success. If children were reasoning about others' suitability as a source of knowledge based on their perceptual access, they were expected to continue to select the knowledgeable informant on switch trials. But, if they were relying on implicit SLSs, they were expected to struggle to continue to be successful in switch trials, because using any associations, or inferred rules, that predicted success in the pre-switch trials would no longer be appropriate.

The results revealed that of those who met criterion, older children were more likely to select the knowledgeable informant in the switch phase trials than younger children. This suggested that older children were more likely to have the capacity to assess the informants' suitability as sources of knowledge. This pattern of results suggested that some of the younger children who had succeeded in the pre-switch trials had done so by forming a simple association, or inferring a rule, based on superficial visual cues. Older children appeared to be more likely to have explicitly reasoned others' perceptual access. They appeared to have used their personal experience of the properties of the frames to infer others' perceptual access and thus assess their suitability as sources of knowledge. This ability was rare in younger children. It may be that younger children did not yet understand that perceptual access grants privileged knowledge, or if they did, they may not have the capacity to appreciate that this renders an individual a valuable source of information. The age-related differences in children's approach to selecting informants was consistent with the expectation that the ability to explicitly reason about others' knowledge develops relatively late in childhood, supporting the view that these capacities are experience-dependent.

Overall, the findings of chapter four were consistent with those in chapter three. They demonstrated that it is possible to differentiate use of implicit and explicitly metacognitive SLSs empirically. Again, the results indicated a relatively late age-related transition from the use of implicit to explicitly metacognitive SLSs.

The Developmental Trajectory of Explicitly Metacognitive SLSs

Across all three studies presented in this thesis children demonstrated age-related improvements in their social information use. These improvements appear to reflect differences in the cognitive capacities that children were employing. Broadly, younger children's performance indicated that they were likely relying on implicit SLSs based on relatively crude and imperfect heuristic biases, associations, or inferred rules. Older children on the other hand, performed in a way that is more similar to the performance we would expect from adults who are assumed to have the capacity for explicitly metacognitive SLSs. Therefore, these results indicate that the capacity for explicitly metacognitive SLSs, in line with predictions, develops relatively late in childhood (Dunstone & Caldwell, 2018; Heyes, 2016c). This late development is also consistent with the hypothesis that explicitly metacognitive SLSs are experience-dependent (Heyes, 2018a). The developmental trend suggests that using explicitly metacognitive SLSs is cognitively challenging.

The age-related changes in social information use, that I have attributed to explicitly metacognitive SLSs, appear to coincide with cognitive developments related to the ability for explicit metacognition. Children's metacognitive ability – their ability to accurately recognise and report their own state of knowledge or ignorance – advances at around 6 years old (S. R. Beck et al., 2012; Kloo et al., 2017; Rohwer et al., 2012; Wimmer et al., 1988). This explicit metacognition is proposed to be unique to humans (Metcalfe, 2015) and is different to the early metacognitive abilities reported in younger children and non-humans (Beran et al., 2015; Goupil et al., 2016; Robinson et al., 2008). The awareness of states of knowledge and ignorance afforded by explicit metacognition might enable reasoning about what information is needed to achieve a particular goal, and who has the potential to provide that information (due to their access to relevant information – through experience or perceptual access). In other words, facilitating recognition of the value of social information. This is further supported by the evidence gained from asking children to verbalise the strategies they were using during the tasks. In chapters three and four children's performance on the task was found to be related to their ability to provide more appropriate verbal reasoning for their choices. That is, children who performed better (more appropriate information seeking and information use) also provided better reasoned responses. This is consistent with the interpretation that some children were indeed employing reasoning-based explicitly metacognitive SLSs based on an understanding of the value of the information.

If the ability to recognise the value, or potential value, of social information is driven by advances in metacognitive ability, then this might also explain why younger children appear to be precluded from recognising the value of others' knowledge. Younger children's apparent reliance on implicit SLSs in the absence of the capacity for explicitly metacognitive SLSs suggests that they are unlikely to have understood the value of the information. Without this understanding, younger children (and by logical extension non-humans) are less likely to use available social information appropriately and effectively.

Having the capacity to recognise the value, to themselves, of social information therefore appears to offer a key advantage in social information use that may drive human cumulative culture beyond the capabilities of non-humans. As such, the characteristic improvement over generations, that we see in adult humans' cumulative culture, could therefore be explained by learners' understanding of the value of the social information that is available from previous generations. If learners understand the value of the information and its relevance to their goal, then it may not matter what form that information comes in. By this I mean that having the ability to recognise the value of social information may make the context in which it is available less important. This may help to explain why the scope of cumulative culture in humans stands in such stark contrast to the limited and context-specific cases reported in non-humans.

Limitations and Future Directions

In this section, I will outline the limitations of the findings reported in this thesis and the scope for further clarification. Of particular interest is the possibility of establishing additional evidence for the implication of specific cognitive developments such as metacognition in children's social information use and whether explicitly metacognitive SLSs are within the capacity of non-humans.

Age as a Proxy for Cognitive Development

All three of the studies presented in this thesis used children's age as a proxy for cognitive development. Therefore, it is not possible to say with certainty whether a particular cognitive development is responsible for, or implicated in, the emergence of the capacity for explicitly metacognitive SLSs. Although I have outlined previously that the findings suggest that advances in metacognitive ability, that occur at around 6 years old, appear to facilitate children's understanding of the relevance and value of social information, it is important to acknowledge that other cognitive developments might be implicated in explicitly metacognitive SLSs. For example, consider advances in executive functions such as inhibitory control or working memory. The development of inhibitory control does suggest that children are reliably able to inhibit a prepotent response by around 6 years old (Gerstadt et al., 1994). In the context of this thesis, if the development of inhibitory control was involved, this may have aided children in their ability to overcome responses related to heuristic biases. However, executive functions tend to develop gradually (Carlson & Moses, 2001). So, if for instance, there had been a more gradual development in appropriate information use in a particular context, it might have indicated that children's performance had improved due to increases in general executive function capabilities. However, the step-changes in performance found at particular ages appears more consistent with the idea that children are making use of a specific cognitive capacity that comes online at that age.

Given that the developmental trajectories of different cognitive capacities overlap it is possible that some of these cognitive developments play a supportive role in the development of metacognitive abilities. Indeed, the metacognitive abilities I refer to above have been proposed to rely upon prior theory of mind understanding (Carruthers, 2009). However, to a certain extent, the specific cognitive development is not the most interesting finding. The key point to note is that the results indicated that younger and older children appeared to tackle the tasks differently, suggesting that there was a fundamental difference in the way they were responding to social information. The developmental trajectory of the age-related changes in children's social information use indicated a relatively late development of the capacity for explicitly metacognitive SLSs. Therefore, even if other cognitive abilities are implicated to some degree in children's ability to reflect on the value of social information, that we do not see advances until relatively late suggests that this capacity is either experience-dependent, or dependent on a relatively mature human cognitive system.

Using age as a proxy for cognitive development was the intended design of these studies. However, more clarity regarding whether particular cognitive developments such as metacognition are implicated could be gained through employing additional control measures such as directly testing children's mental state understanding and metacognitive abilities. For example, alongside tasks such as those in the preceding chapters, participants could also undergo a battery of standardised cognitive tests. If performance on such tests is correlated with appropriate responses in the social information use tasks, it could bolster the argument for the involvement of that capacity, though it would not confirm whether that capacity is a necessary prerequisite. Another alternative way to establish the implication of particular cognitive capacities in distinctively human cumulative culture would be to restrict access to capacities of interest in human adults (Dunstone & Caldwell, 2018). For example, assessing the role of explicit metacognition in the kind of flexible social learning that is presumed necessary for distinctively human cumulative culture requires the availability of metacognitive capacities to be manipulated. However, the capacity for explicit metacognition cannot be removed experimentally, therefore, one way to assess its involvement could be restricting access to the explicit cognitive processes it depends on. In particular, given that executive functions are proposed to rely on similar explicit cognitive processes, employing dual task methods that tax executive functions could act as a proxy for restricting explicit metacognition (see approach taken by Dunstone et al., 2021). If explicit cognitive processes, such as explicit metacognition, are involved in the kind of social learning required for the emergence of cumulative culture, the expectation would be that restricting their availability would reduce the capacity for flexible social learning decisions.

A Case of Indirect Evidence

This thesis has repeatedly claimed that evidence of the late development of explicitly metacognitive SLSs in children suggests that they are likely to be beyond the capabilities of non-humans. This conclusion seems plausible and is consistent with the fact that much of the literature on SLSs in non-humans can be retrospectively attributed to implicit SLSs (Heyes, 2017; Heyes & Pearce, 2015). However, some researchers may argue that the lack of direct evidence that non-humans cannot use explicitly metacognitive SLSs is insufficient for making this claim. Indeed, the conclusions drawn in this thesis in relation to the distinctiveness of human cumulative culture are naturally based to some extent on conjecture. Therefore, to support the proposal that young children and non-humans rely on implicit SLSs in circumstances in which the ability to employ explicitly metacognitive SLSs would be advantageous, it might be beneficial to assess non-humans directly. If non-humans were to perform in a manner consistent with the pattern of results we see in younger children then this would offer further credibility to the claims made in this thesis, as well as to the proposal that explicitly metacognitive SLSs are implicated in distinctively human cumulative culture. However, such a study would likely be difficult to design and extremely challenging to implement.

To date, there have been no plausible solutions offered with regard to assessing explicit knowledge, relative to implicit knowledge, in non-humans. Indeed, this is a persistent problem in the study of cognition in non-humans (Gómez et al., 2017). In particular, these difficulties relate to problems with testing explicit understanding in non-verbal species. In the absence of opportunities to explain to subjects the aim and structure of a task, or to ask them directly about their reasoning or their understanding of a situation, we cannot be certain about what they might have inferred about the challenge presented to them, or what they are aiming to achieve. While extensive training can help to establish the goals of a task it considerably restricts any interpretations of subsequent performance, inasmuch as it can usually be attributed to associations, or rule-like strategies, formed during training. Therefore, the certainty with which it would be possible to say whether or not subjects have an explicit understanding would be extremely limited. Added to this, claims of metacognition in non-humans remain contentious (Carruthers, 2008; Carruthers & Ritchie, 2012). For example, Carruthers (2008) argued that the purported metacognitive responses in non-humans could be explainable by simpler cognitive processes such as first-order beliefs or deriving a general rule from training conditions. The challenges associated with assessing explicit knowledge (and metacognition more broadly) present a considerable problem for the hypothesis that non-humans do not have the capacity for explicitly metacognitive SLSs. That is, without a suitable way to test these capacities in non-humans this hypothesis is currently unfalsifiable. However, such obstacles to testing this hypothesis in non-human populations reinforces the importance of the more indirect approach taken in this thesis. As such, looking to determine the age at which children develop the capacity for explicitly metacognitive SLSs, which cognitive mechanisms are likely involved in its development, and whether there is evidence of such cognitive mechanisms in non-humans, offers an indirect way to assess the likelihood of non-humans having the capacity for explicitly metacognitive SLSs.

Concluding Remarks

This thesis examined age-related changes in children's social information use in response to particular cognitive challenges as a means to understand the cognitive processes that might be implicated in distinctively human cumulative culture. Overall, the studies that I have presented in this thesis provide support for the proposal that the kind of social information use that affords distinctively human cumulative culture is facilitated by explicitly metacognitive SLSs. These reasoning-based strategies appear to be driven by an ability to understand the value, or potential value, of social information. That is, the flexibility afforded by the ability to recognise its value appears to offer the significant advantage in social information use that may drive human cumulative culture beyond the capabilities of non-humans. Indications that the development of explicitly metacognitive SLSs occurs relatively late in childhood offers credence to the assertion that this capacity is likely to be experience-dependent and unlikely to be observed in non-humans. Specifically, it appears that experience-dependent metacognitive abilities and mental state understanding have a relevant influence on social information use. Therefore, if the real world contexts in which we see cumulative culture in humans involve cognitive challenges such as those investigated in this thesis, this might explain why there is limited evidence of cumulative culture in non-humans.

References

Alem, S., Perry, C. J., Zhu, X., Loukola, O. J., Ingraham, T., Søvik, E., & Chittka, L. (2016). Associative mechanisms allow for social learning and cultural transmission of string pulling in an insect. *PLoS Biology*, *14*(10), 1–28. https://doi.org/10/f89jjs

Allen, J. (2019). Community through culture: From insects to whales: How social learning and culture manifest across diverse animal communities. *BioEssays*, *41*(11), 1–8. https://doi.org/10/ggb5kq

Allen, J., Weinrich, M., Hoppitt, W., & Rendell, L. (2013). Network-based diffusion analysis reveals cultural transmission of lobtail feeding in humpback whales. *Science*, *340*(6131), 485–488. https://doi.org/10/f4vbds

Anderson, P. (2002). Assessment and development of executive function (EF) during childhood. *Child Neuropsychology*, *8*(2), 71–82. https://doi.org/10/bn6vj8

Anderson, P., Anderson, V., Northam, E., & Taylor, H. G. (2000). Standardization of the Contingency Naming Test (CNT) for school-age children: A measure of reactive flexibility. *Clinical Neuropsychological Assessment*, *1*, 247–273.

Aplin, L. M. (2019). Culture and cultural evolution in birds: A review of the evidence. *Animal Behaviour*, *147*, 179–187. https://doi.org/10/gfsp4p

Aplin, L. M., Farine, D. R., Morand-Ferron, J., Cockburn, A., Thornton, A., & Sheldon, B. C. (2015). Experimentally induced innovations lead to persistent culture via conformity in wild birds. *Nature*, *518*(7540), 538–541. https://doi.org/10/f3pfvt

Arbesman, S. (2011). Quantifying the ease of scientific discovery. *Scientometrics*, *86*(2), 245–250. https://doi.org/10/fgdrh3

Atkinson, M., Renner, E., Thompson, B., Mackintosh, G., Xie, D., Su, Y., & Caldwell, C. A. (2020). Robust, source-independent biases in children's use of socially and individually acquired information. *Journal of Experimental Psychology: General*. https://doi.org/10/fmft

Baddeley, A. D., & Hitch, G. (1974). Working memory. *Psychology of Learning and*

Motivation - Advances in Research and Theory, *8*(C), 47–89. https://doi.org/10/d4vgnk

Baillargeon, R., Scott, R. M., & He, Z. (2010). False-belief understanding in infants. *Trends in Cognitive Sciences*, *14*(3), 110–118. https://doi.org/10/c9qntv

Baillargeon, R., Spelke, E. S., & Wasserman, S. (1985). Object permanence in five-month-old infants. *Cognition*, *20*(3), 191–208. https://doi.org/10/cvxsbc

Balcomb, F. K., & Gerken, L. A. (2008). Three-year-old children can access their own memory to guide responses on a visual matching task. *Developmental Science*, *11*(5), 750–760. https://doi.org/10/dj6459

Baldwin, D. A., & Moses, L. J. (1996). The ontogeny of social information gathering. *Child Development*, *67*(5), 1915–1939. https://doi.org/10/cjs68g

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 1–43. https://doi.org/10/gcm4wc

Barrouillet, P. (2015). Theories of cognitive development: From Piaget to today. *Developmental Review*, *38*, 1–12. https://doi.org/10/gddnvj

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting Linear Mixed-Effects Models using lme4. *Journal of Statistical Software*, *67*(1). https://doi.org/10/gcrnkw

Beck, S. R., McColgan, K. L. T., Robinson, E. J., & Rowley, M. G. (2011). Imagining what might be: Why children underestimate uncertainty. *Journal of Experimental Child Psychology*, *110*(4), 603–610. https://doi.org/10/fq2srg

Beck, S. R., Robinson, E. J., & Rowley, M. G. (2012). Thinking about different types of uncertainty. In M. J. Beran, J. Brandi, & J. Proust (Eds.), *Foundations of Metacognition* (pp. 181–192). Oxford University Press. https://doi.org/10/fmfx

Beck, W. (1992). Aboriginal preparation of cycas seeds in Australia. *Economic Botany*, *46*(2), 133–147.

Beran, M. J., Perdue, B. M., Futch, S. E., Smith, J. D., Evans, T. A., & Parrish, A. E. (2015). Go when you know: Chimpanzees' confidence movements reflect their responses in a computerized memory task. *Cognition*, *142*, 236–246. https://doi.org/10/f7mdzq

Bernard, S., Proust, J., & Clément, F. (2015). Procedural metacognition and false belief understanding in 3- to 5-year-old children. *PLoS ONE*, *10*(10), 1–12. https://doi.org/10/drb4

Best, J. R., & Miller, P. H. (2010). A developmental perspective on executive function. *Child Development*, *81*(6), 1641–1660. https://doi.org/10/fmbqgm

Blakey, K. H., Rafetseder, E., Atkinson, M., Renner, E., Cowan-Forsythe, F., Sati, S. J., & Caldwell, C. A. (2020). Development of strategic social information seeking: Implications for cumulative culture. *PsyArXiv*. https://doi.org/10.31234/osf.io/4envx

Blanco, N. J., & Sloutsky, V. M. (2020). Systematic exploration and uncertainty dominate young children's choices. *Developmental Science*. https://doi.org/10/gg7whd

Bloom, P., & Markson, L. (1998). Intention and analogy in children's naming of pictorial representations. *Psychological Science*, *9*(3), 200–204. https://doi.org/10/c9hxz2

Boyd, R., & Richerson, P. J. (1996). Why culture is common, but cultural evolution is rare. *Proceedings of the British Academy*, *88*, 77–93.

Brosseau-Liard, P. E., & Birch, S. A. J. (2011). Epistemic states and traits: Preschoolers appreciate the differential informativeness of situation-specific and person-specific cues to knowledge. *Child Development*, *82*(6), 1788–1796. https://doi.org/10/dgzstg

Brydges, C. R., Fox, A. M., Reid, C. L., & Anderson, M. (2014). The differentiation of executive functions in middle and late childhood: A longitudinal latent-variable analysis. *Intelligence*, *47*, 34–43. https://doi.org/10/f6sm2j

Burdett, E. R. R., Lucas, A. J., Buchsbaum, D., McGuigan, N., Wood, L. A., & Whiten, A. (2016). Do children copy an expert or a majority? Examining selective learning in instrumental and normative contexts. *PLoS ONE*, *11*(10), 15–17. https://doi.org/10/f9rmhz

Butler, L. P., Gibbs, H. M., & Tavassolie, N. S. (2020). Children's developing understanding that even reliable sources need to verify their claims. *Cognitive Development*, *54*, 100871. https://doi.org/10/dtsk

Buttelmann, D., Carpenter, M., & Tomasello, M. (2009). Eighteen-month-old infants show false belief understanding in an active helping paradigm. *Cognition*, *112*(2), 337–342. https://doi.org/10/cv6x5b

Caldwell, C. A. (2018). Using experimental research designs to explore the scope of cumulative culture in humans and other animals. *Topics in Cognitive Science*, 1–17. https://doi.org/10/djsh

Caldwell, C. A., Atkinson, M., Blakey, K. H., Dunstone, J., Kean, D., Mackintosh, G., Renner, E., & Wilks, C. E. H. (2020). Experimental assessment of capacities for cumulative culture: Review and evaluation of methods. *Wiley Interdisciplinary Reviews: Cognitive Science*, *11*(1), 1–15. https://doi.org/10/djm2

Caldwell, C. A., Atkinson, M., & Renner, E. (2016). Experimental approaches to studying cumulative cultural evolution. *Current Directions in Psychological Science*, *25*(3), 191–195. https://doi.org/10/f8tgzx

Caldwell, C. A., & Millen, A. E. (2008a). Experimental models for testing hypotheses about cumulative cultural evolution. *Evolution and Human Behavior*, *29*(3), 165–171. https://doi.org/10/fhqgrd

Caldwell, C. A., & Millen, A. E. (2008b). Studying cumulative cultural evolution in the laboratory. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *363*(1509), 3529–3539. https://doi.org/10/ckpzmq

Caldwell, C. A., & Millen, A. E. (2009). Social learning mechanisms and cumulative cultural evolution: Is imitation necessary? *Psychological Science*, *20*(12), 1478–1483. https://doi.org/10/fs6nx7

Caldwell, C. A., Renner, E., & Atkinson, M. (2018). Human teaching and cumulative cultural evolution. *Review of Philosophy and Psychology*, *9*(4), 751–770. https://doi.org/10/gf8hdw

Caldwell, C. A., Schillinger, K., Evans, C. L., & Hopper, L. M. (2012). End state copying by humans (*Homo sapiens*): Implications for a comparative perspective on cumulative culture. *Journal of Comparative Psychology*, *126*(2), 161–169. https://doi.org/10/f3zx24

Call, J., & Carpenter, M. (2001). Do apes and children know what they have seen? *Animal Cognition*, *3*(4), 207–220. https://doi.org/10/dw8p4p

Call, J., & Tomasello, M. (1999). A non-verbal false belief task: the performance of chimpanzees and human children. *Child Development*, *70*(2), 381–395. https://doi.org/10/cp3kdt

Carlson, S. M., & Moses, L. J. (2001). Individual differences in inhibitory control and children's theory of mind. *Child Development*, *72*(4), 1032–1053. https://doi.org/10/dfpng3

Carr, K., Kendal, R., & Flynn, E. G. (2015). Imitate or innovate? Children's innovation is influenced by the efficacy of observed behaviour. *Cognition*, *142*, 322–332. https://doi.org/10/f7md9h

Carruthers, P. (2008). Meta-cognition in animals: A skeptical look. *Mind and Language*, *23*(1), 58–89. https://doi.org/10/c935m7

Carruthers, P. (2009). How we know our own minds: The relationship between mindreading and metacognition. *Behavioral and Brain Sciences*, *32*(2), 121–182. https://doi.org/10/fd7dbt

Carruthers, P., & Ritchie, J. B. (2012). The emergence of metacognition: Affect and uncertainty in animals. In *Foundations of Metacognition* (pp. 76–93). https://doi.org/10/fmfw

Christensen, R. H. B. (2019). *Ordinal - Regression models for ordinal data*. https://cran.r-project.org/package=ordinal

Corriveau, K. H., & Harris, P. L. (2009). Choosing your informant: Weighing familiarity and recent accuracy. *Developmental Science*, *12*(3), 426–437. https://doi.org/10/bwssfs

Corriveau, K. H., Kinzler, K. D., & Harris, P. L. (2013). Accuracy trumps accent in children's endorsement of object labels. *Developmental Psychology*, *49*(3), 470–479. https://doi.org/10/f4pwg8

Danchin, E., Nöbel, S., Pocheville, A., Dagaeff, A.-C., Demay, L., Alphand, M., Ranty-Roby, S., van Renssen, L., Monier, M., Gazagne, E., Allain, M., & Isabel, G. (2018). Conformist social learning in fruitflies predicts long-lasting mate-choice traditions. *Science*, *362*, 1025–1030. https://doi.org/10/gfq734

Davidson, M. C., Amso, D., Anderson, L. C., & Diamond, A. (2006). Development of cognitive control and executive functions from 4 to 13 years: Evidence from manipulations of memory, inhibition, and task switching. *Neuropsychologia*, *44*(11), 2037–2078. https://doi.org/10/bfs49w

De Oliveira, E., Reynaud, E., & Osiurak, F. (2019). Distinct roles of theory of mind and technical reasoning in cumulative technological culture. *Human Nature*, *30*, 326–340. https://doi.org/10/djmz

Dean, L. G., Kendal, R. L., Schapiro, S. J., Thierry, B., & Laland, K. N. (2012). Identification of the social and cognitive processes underlying human cumulative culture. *Science*, *335*(6072), 1114–1118. https://doi.org/10/gc3j85

Dean, L. G., Vale, G., Laland, K. N., Flynn, E. G., & Kendal, R. (2014). Human cumulative culture: A comparative perspective. *Biological Reviews*, *89*(2), 284–301. https://doi.org/10/f5356m

Dennett, D. C. (1978). Beliefs about beliefs [P&W, SR&B]. *Behavioral and Brain Sciences*, *1*(4), 568–570. https://doi.org/10/cfzp8g

Derex, M., Bonnefon, J. F., Boyd, R., & Mesoudi, A. (2019). Causal understanding is not necessary for the improvement of culturally evolving technology. *Nature Human Behaviour*, *3*(5), 446–452. https://doi.org/10/gghxhk

Derex, M., Perreault, C., & Boyd, R. (2018). Divide and conquer: Intermediate levels of population fragmentation maximize cultural accumulation. *Philosophical Transactions of the Royal Society B: Biological Sciences, 373*(1743). https://doi.org/10/fmfh

Diamond, A. (2006). The early development of executive functions. In E. Bialystok & F. I. M. Craik (Eds.), *Lifespan Cognition: Mechanisms of Change* (pp. 70–95). Oxford University Press. https://doi.org/10/gf66qq

Diamond, A. (2013). Executive functions. *Annual Review of Psychology*, *64*, 135–168. https://doi.org/10/b2m2

Dunham, P., Dunham, F., & O'Keefe, C. (2000). Two-year-olds' sensitivity to a parent's knowledge state: Mind reading or contextual cues? *British Journal of Developmental Psychology*, *18*(4), 519–532. https://doi.org/10/drg3ft

Dunstone, J., Atkinson, M., Grainger, C., Renner, E., & Caldwell, C. A. (2021). Limited evidence for executive function load impairing selective copying in a win-stay lose-shift task. *PLoS ONE*, *16*(3), e0247183. https://doi.org/10/fx8r

Dunstone, J., & Caldwell, C. A. (2018). Cumulative culture and explicit metacognition: A review of theories, evidence and key predictions. *Palgrave Communications*, *4*(1), 145. https://doi.org/10/djm3

Enquist, M., Eriksson, K., & Ghirlanda, S. (2007). Critical social learning: A solution to Rogers's paradox of nonadaptive culture. *American Anthropologist*, *109*(4), 727–734. https://doi.org/10/ddpfb2

Enquist, M., Ghirlanda, S., & Eriksson, K. (2011). Modelling the evolution and diversity of cumulative culture. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1563), 412–423. https://doi.org/10/bnqfct

Enquist, M., Strimling, P., Eriksson, K., Laland, K., & Sjostrand, J. (2010). One cultural parent makes no culture. *Animal Behaviour*, *79*(6), 1353–1362. https://doi.org/10/b9zwhp

Evans, C. L., Laland, K. N., Carpenter, M., & Kendal, R. (2018). Selective copying of the majority suggests children are broadly "optimal-" rather than "over-" imitators. *Developmental Science*, *21*(5), 1–10. https://doi.org/10/gcpj8w

Fay, N., De Kleine, N., Walker, B., & Caldwell, C. A. (2019). Increasing population size can inhibit cumulative cultural evolution. *Proceedings of the National Academy of Sciences*, *116*(14), 6726–6731. https://doi.org/10/ggb8nj

Fedra, E., & Schmidt, M. F. H. (2019). Older (but not younger) preschoolers reject incorrect knowledge claims. *British Journal of Developmental Psychology*, *37*(1), 130–145. https://doi.org/10/djs8

Flavell, J. H. (1971). Stage-related properties of cognitive development. *Cognitive Psychology*, *2*(4), 421–453. https://doi.org/10/fs7qj8

Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive — developmental inquiry. *American Psychologist*, *34*(10), 906–911. https://doi.org/10/cg33qs

Flavell, J. H., Mumme, D. L., Green, F. L., & Flavell, E. R. (1992). Young children's understanding of different types of beliefs. *Child Development*, *63*(4), 960–977. https://doi.org/10/bp3hsb

Flynn, E. G. (2008). Investigating children as cultural magnets: Do young children transmit redundant information along diffusion chains? *Philosophical Transactions of the Royal Society B: Biological Sciences*, *363*(1509), 3541–3551. https://doi.org/10/d52g48

Flynn, E. G., & Whiten, A. (2013). Dissecting children's observational learning of complex actions through selective video displays. *Journal of Experimental Child Psychology*, *116*(2), 247–263. https://doi.org/10/f5fhjb

Foote, A. L., & Crystal, J. D. (2007). Metacognition in the rat. *Current Biology*, *17*(6), 551–555. https://doi.org/10/bmd8ts

Garland, E. C., Goldizen, A. W., Rekdahl, M. L., Constantine, R., Garrigue, C., Hauser, N. D., Poole, M. M., Robbins, J., & Noad, M. J. (2011). Dynamic horizontal cultural transmission of humpback whale song at the ocean basin scale. *Current*

Biology, *21*(8), 687–691. https://doi.org/10/bf8bb2

Garland, E. C., & McGregor, P. K. (2020). Cultural transmission, evolution, and revolution in vocal displays: Insights from bird and whale song. *Frontiers in Psychology*, *11*, 544929. https://doi.org/10/fmfk

Gathercole, S. E., Pickering, S. J., Ambridge, B., & Wearing, H. (2004). The structure of working memory from 4 to 15 years of age. *Developmental Psychology*, *40*(2), 177–190. https://doi.org/10/bc24qv

Gerstadt, C. L., Hong, Y. J., & Diamond, A. (1994). The relationship between cognition and action: Performance of children 3.5 - 7 years old on a Stroop-like day-night test. *Cognition*, *53,* 129–153. https://doi.org/10/bzvbzx

Gómez, J.-C., Kersken, V., Ball, D., & Seed, A. (2017). Knowing without knowing: Implicit cognition and the minds of infants and animals / Saber sin saber: la cognición implícita y las mentes de niños pequeños y animales. *Estudios de Psicologia*, *38*(1), 37–62. https://doi.org/10/fmfv

Gopnik, A. (2020). Childhood as a solution to explore-exploit tensions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *375*(1803), 20190502. https://doi.org/10/ggzkn5

Gopnik, A., & Astington, J. W. (1988). Children's understanding of representational change and its relation to the understanding of false belief and the appearance-reality distinction. *Child Development*, *59*(1), 26. https://doi.org/10/d9gp6v

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. https://doi.org/10/f8gq87

Harris, P. L., Koenig, M. A., Corriveau, K. H., & Jaswal, V. K. (2018). Cognitive foundations of learning from testimony. *Annual Review of Psychology*, *69*(1), 251–273. https://doi.org/10/gdrbzv

Haselgrove, M. (2016). Overcoming associative learning. *Journal of Comparative Psychology*, *130*(3), 226–240. https://doi.org/10/f84dcx

Haun, D. B. M., Rekers, Y., & Tomasello, M. (2012). Majority-biased transmission in chimpanzees and human children, but not orangutans. *Current Biology*, *22*(8), 727–731. https://doi.org/10/f3xn46

Haun, D. B. M., van Leeuwen, E. J. C., & Edelson, M. G. (2013). Majority influence in children and other animals. *Developmental Cognitive Neuroscience*, *3*, 61–71. https://doi.org/10/djsk

Hawthorne-Madell, D., & Goodman, N. D. (2019). Reasoning about social sources to learn from actions and outcomes. *Decision*, *6*(1), 17–60. https://doi.org/10/djtc

Henrich, J. (2016). *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. Princeton University Press. https://doi.org/10/fmfm

Henrich, J., & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology*, *12*(3), 123–135. https://doi.org/10/bg7bf5

Herschensohn, J. (2007). Right on time: process and schedule of first language acquisition. In *Language Development and Age* (pp. 27–64). Cambridge University Press. https://doi.org/10/cjq49p

Heyes, C. (1998). Theory of mind in nonhuman primates. *Behavioral and Brain Sciences*, *21*(1), 101–148. https://doi.org/10/bs657r

Heyes, C. (2011). Automatic imitation. *Psychological Bulletin*, *137*(3), 463–483. https://doi.org/10/ffd5r8

Heyes, C. (2012). Grist and mills: On the cultural origins of cultural learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *367*(1599), 2181–2191. https://doi.org/10/f36dqx

Heyes, C. (2016a). Blackboxing: Social learning strategies and cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *371*(1693). https://doi.org/10/gfzkn4

Heyes, C. (2016b). Born pupils? Natural pedagogy and cultural pedagogy. *Perspectives on Psychological Science*, *11*(2), 280–295. https://doi.org/10/f8f6m3

Heyes, C. (2016c). Who knows? Metacognitive social learning strategies. *Trends in Cognitive Sciences*, *20*(3), 204–213. https://doi.org/10/f8bpsm

Heyes, C. (2017). When does social learning become cultural learning? *Developmental Science*, *20*(2), 1–14. https://doi.org/10/gfzkt4

Heyes, C. (2018a). *Cognitive gadgets: The cultural evolution of thinking*. The Belknap Press of Harvard University Press.

Heyes, C. (2018b). Enquire within: Cultural evolution and cognitive science. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1743), 1–9. https://doi.org/10/djpz

Heyes, C. (2018c). Précis of cognitive gadgets: The cultural evolution of thinking. *Behavioral and Brain Sciences*, *42*, 1–58. https://doi.org/10/gfv325

Heyes, C. (2019). Is morality a gadget? Nature, nurture and culture in moral development. *Synthese*. https://doi.org/10/ghcs74

Heyes, C., & Pearce, J. M. (2015). Not-so-social learning strategies. *Proc. R. Soc. B*, *282*. https://doi.org/10/gfzkr8

Hobaiter, C., Poisot, T., Zuberbühler, K., Hoppitt, W., & Gruber, T. (2014). Social network analysis shows direct evidence for social transmission of tool use in wild chimpanzees. *PLoS Biology*, *12*(9). https://doi.org/10/f6j6zv

Hoehl, S., Keupp, S., Schleihauf, H., McGuigan, N., Buttelmann, D., & Whiten, A. (2019). ‘Over-imitation’: A review and appraisal of a decade of research. *Developmental Review*, *51*, 90–108. https://doi.org/10/gfspt3

Hoppitt, W. J. E., & Laland, K. N. (2013). *Social learning: An introduction to mechanisms, methods, and models*. Princeton University Press.

Horner, V., Proctor, D., Bonnie, K. E., Whiten, A., & de Waal, F. B. M. (2010). Prestige affects cultural learning in chimpanzees. *PLoS ONE*, *5*(5), 1–5. https://doi.org/10/drfvhw

Horner, V., & Whiten, A. (2007). Learning from others’ mistakes? Limits on understanding a trap-tube task by young chimpanzees (*Pan troglodytes*) and children (*Homo sapiens*). *Journal of Comparative Psychology*, *121*(1), 12–21. https://doi.org/10/fm362w

Huber, L., Range, F., Voelkl, B., Szucsich, A., Virányi, Z., & Miklosi, A. (2009). The evolution of imitation: What do the capacities of non-human animals tell us about the mechanisms of imitation? *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1528), 2299–2309. https://doi.org/10/dbvrnz

Hughes, C. (1998a). Executive function in preschoolers: Links with theory of mind and verbal ability. *British Journal of Developmental Psychology*, *16*, 233–253. https://doi.org/10/dn5d58

Hughes, C. (1998b). Finding your marbles: Does preschoolers' strategic behavior predict later understanding of mind? *Developmental Psychology*, *34*(6), 1326–1339. https://doi.org/10/dbzg9g

Hunt, G. R., & Gray, R. D. (2003). Diversification and cumulative evolution in New Caledonian crow tool manufacture. *Proceedings of the Royal Society B: Biological Sciences*, *270*(1517), 867–874. https://doi.org/10/b4wm8v

Jaswal, V. K., & Neely, L. A. (2006). Adults don't always know best: Preschoolers use past reliability over age when learning new words. *Psychological Science*, *17*(9), 757–758. https://doi.org/10/bsf845

Jenkins, J. M., & Astington, J. W. (1996). Cognitive factors and family structure associated with theory of mind development in youong children. *Developmental Psychology*, *32*(1), 70–78.

Jesmer, B. R., Merkle, J. A., Goheen, J. R., Aikens, E. O., Beck, J. L., Courtemanch, A. B., Hurley, M. A., McWhirter, D. E., Miyasaki, H. M., Monteith, K. L., & Kauffman, M. J. (2018). Is ungulate migration culturally transmitted? Evidence of social learning from translocated animals. *Science*, *361*(6406), 1023–1025. https://doi.org/10/ctpm

Jiménez, Á. V., & Mesoudi, A. (2019). Prestige-biased social learning: Current evidence and outstanding questions. *Palgrave Communications*, *5*(20), 1–12.

https://doi.org/10/gf9hv2

Jin, X., Li, P., He, J., & Shen, M. (2017). Cooperation, but not competition, improves 4-year-old children's reasoning about others' diverse desires. *Journal of Experimental Child Psychology*, *157*, 81–94. https://doi.org/10/f9zp33

Kano, F., Krupenye, C., Hirata, S., Tomonaga, M., & Call, J. (2019). Great apes use self-experience to anticipate an agent's action in a false-belief test. *Proceedings of the National Academy of Sciences*, 1–6. https://doi.org/10/gf85q5

Karg, K., Schmelz, M., Call, J., & Tomasello, M. (2015). The goggles experiment: Can chimpanzees use self-experience to infer what a competitor can see? *Animal Behaviour*, *105*, 211–221. https://doi.org/10/djtd

Kempe, M., Lycett, S. J., & Mesoudi, A. (2014). From cultural traditions to cumulative culture: Parameterizing the differences between human and nonhuman culture. *Journal of Theoretical Biology*, *359*, 29–36. https://doi.org/10/f6fzng

Kendal, R., Boogert, N. J., Rendell, L., Laland, K. N., Webster, M. M., & Jones, P. L. (2018). Social learning strategies: Bridge-building between fields. *Trends in Cognitive Sciences*, *22*(7), 651–665. https://doi.org/10/gds7v2

Kendal, R., Hopper, L. M., Whiten, A., Brosnan, S. F., Lambeth, S. P., Schapiro, S. J., & Hoppitt, W. (2015). Chimpanzees copy dominant and knowledgeable individuals: Implications for cultural diversity. *Evolution and Human Behavior*, *36*(1), 65–72. https://doi.org/10/gfsvv2

Kloo, D., & Rohwer, M. (2012). The development of earlier and later forms of metacognitive abilities: Reflections on agency and ignorance. In M. J. Beran, J. Brandi, & J. Proust (Eds.), *Foundations of Metacognition* (pp. 583–605). Oxford University Press. https://doi.org/10/dpmg

Kloo, D., Rohwer, M., & Perner, J. (2017). Direct and indirect admission of ignorance by children. *Journal of Experimental Child Psychology*, *159*, 279–295. https://doi.org/10/f97m77

Kloo, D., & Sodian, B. (2017). The developmental stability of inhibition from 2 to 5 years. *British Journal of Developmental Psychology*, *35*(4), 582–595. https://doi.org/10/gcc2c9

Koenig, M. A., Clément, F., & Harris, P. L. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science*, *15*(10), 694–698. https://doi.org/10/ft5xqk

Koenig, M. A., & Harris, P. L. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, *76*(6), 1261–1277. https://doi.org/10/dw4t27

Koenig, M. A., Tiberius, V., & Hamlin, J. K. (2019). Children's judgments of epistemic and moral agents: From situations to intentions. *Perspectives on Psychological Science*, *14*(3), 344–360. https://doi.org/10/drb5

Kovács, Á. M., Téglás, E., & Endress, A. D. (2010). The social sense: Susceptibility to others' beliefs in human infants and adults. *Science*, *330*, 1830–1834. https://doi.org/10/cq6gmc

Krachun, C., Carpenter, M., Call, J., & Tomasello, M. (2009). A competitive nonverbal false belief task for children and apes. *Developmental Science*, *12*(4), 521–535. https://doi.org/10/fvccj2

Krupenye, C., & Call, J. (2019). Theory of mind in animals: Current and future directions. *Wiley Interdisciplinary Reviews: Cognitive Science*, *April*, 1–25. https://doi.org/10/djtf

Krupenye, C., Kano, F., Hirata, S., Call, J., & Tomasello, M. (2016). Great apes anticipate that other individuals will act according to false beliefs. *Science*, *354*(6308), 110–114. https://doi.org/10/brmq

Kuhn, D. (2000). Metacognitive development. *Current Directions in Psychological Science*, *9*(5), 178–181. https://doi.org/10/cnp5fp

Kulke, L., Johannsen, J., & Rakoczy, H. (2019). Why can some implicit theory of mind tasks be replicated and others cannot? A test of mentalizing versus submentalizing accounts. *PLoS ONE*, *14*(3). https://doi.org/10/ggb8pv

Laland, K. N. (2004). Social learning strategies. *Animal Learning & Behavior*, *32*(1), 4–

14. https://doi.org/10/fbt28j

Lenth, R., Singmann, H., Love, J., Buerkner, P., & Herve, M. (2019). *Emmeans - Estimated marginal means, aka least-squares means*. https://cran.r-project.org/package=emmeans

Lewis, H. M., & Laland, K. N. (2012). Transmission fidelity is the key to the build-up of cumulative culture. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *367*(1599), 2171–2180. https://doi.org/10/f3573w

Lucas, A. J., Burdett, E. R. R., Burgess, V., Wood, L. A., McGuigan, N., Harris, P. L., & Whiten, A. (2017). The development of selective copying: Children's learning from an expert versus their mother. *Child Development*, *88*(6), 2026–2042. https://doi.org/10/gckt88

Lundberg, I. (2010). Early language development as related to the acquisition of reading. *Cognition and Neuropsychology: International Perspectives on Psychological Science*, *14*(1), 193–208. https://doi.org/10/fm6r

Lurz, R., & Krachun, C. (2019). Experience-projection methods in theory-of-mind research: Their limits and strengths. *Current Directions in Psychological Science*, *28*(5), 456–462. https://doi.org/10/drb6

May, K. O. (1966). Quantitative growth of the mathematical literature. *Science*, *154*, 1672–1673.

McGarrigle, J., & Donaldson, M. (1974). Conservation accidents. *Cognition*, *3*(4), 341–350. https://doi.org/10/cm3dgg

McGuigan, N. (2013). The influence of model status on the tendency of young children to over-imitate. *Journal of Experimental Child Psychology*, *116*(4), 962–969. https://doi.org/10/f5hqn6

McGuigan, N., Burdett, E. R. R., Burgess, V., Dean, L. G., Lucas, A. J., Vale, G., & Whiten, A. (2017). Innovation and social transmission in experimental micro-societies: Exploring the scope of cumulative culture in young children. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*(1735), 20160425. https://doi.org/10/fmfn

Meltzoff, A. N., & Brooks, R. (2008). Self-experience as a mechanism for learning about others: A training study in social cognition. *Developmental Psychology*, *44*(5), 1257–1265. https://doi.org/10/fcfk7v

Mesoudi, A., & Thornton, A. (2018). What is cumulative cultural evolution? *Proc. R. Soc. B*, *285*. https://doi.org/10/gf69rw

Metcalfe, J. (2015). Evolution of metacognition. *Handbook of Metamemory and Memory*, *11010*. https://doi.org/10/fmfp

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*(May), 167–202. https://doi.org/10/fhpgvb

Mills, C. M. (2013). Knowing when to doubt: Developing a critical stance when learning from others. *Developmental Psychology*, *49*(3). https://doi.org/10/f4p544

Miu, E., Gulley, N., Laland, K. N., & Rendell, L. (2020). Flexible learning, rather than inveterate innovation or copying, drives cumulative knowledge gain. *Science Advances*, *6*(6), 1–10. https://doi.org/10/gg3dvj

Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wager, T. D. (2000). The unity and diversity of executive functions and their contributions to complex “frontal lobe” tasks: A latent variable analysis. *Cognitive Psychology*, *41*(1), 49–100. https://doi.org/10/bkksp2

Moore, C., Jarrold, C., Russell, J., Lumb, A., Sapp, F., & MacCalIum, F. (1995). Conflicting desire and the child’s theory of mind. *Cognitive Development*, *10*(4), 467–482. https://doi.org/10/d4z8r4

Morgan, T. J. H., Rendell, L. E., Ehn, M., Hoppitt, W., & Laland, K. N. (2012). The evolutionary basis of human social learning. *Proc. R. Soc. B*, *279*(1729), 653–662. https://doi.org/10/ccjjb8

Moses, L. J., Coon, J. A., & Wusinich, N. (2000). Young children’s understanding of desire formation. *Developmental Psychology*, *36*(1), 77–90. https://doi.org/10/bnf6t8

Muthukrishna, M., Shulman, B. W., Vasilescu, V., & Henrich, J. (2014). Sociality influences cultural complexity. *Proceedings of the Royal Society B: Biological Sciences*, *281*, 20132511. https://doi.org/10/gfzksp

Neldner, K., Collier-Baker, E., & Nielsen, M. (2015). Chimpanzees (*Pan troglodytes*) and human children (*Homo sapiens*) know when they are ignorant about the location of food. *Animal Cognition*, *18*(3), 683–699. https://doi.org/10/f68jr7

Nickerson, R. S. (1999). How we know - and sometimes misjudge - what others know: Imputing one's own knowledge to others. *Psychological Bulletin*, *125*(6), 737–759. https://doi.org/10/fmh82f

O'Neill, D. K. (1996). Two-year-old children's sensitivity to a parent's knowledge state when making requests. *Child Development*, *67*(2), 659–677. https://doi.org/10/cqqnks

O'Neill, D. K., Astington, J. W., & Flavell, J. H. (1992). Young children's understanding of the role that sensory experiences play in knowledge acquisition. *Child Development*, *63*(2), 474–490. https://doi.org/10/dtrv3p

Onishi, K. H., & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, *308*(5719), 255–258. https://doi.org/10/bj7ndt

Ottoni, E. B., De Resende, B. D., & Izar, P. (2005). Watching the best nutcrackers: What capuchin monkeys (*Cebus apella*) know about others' tool-using skills. *Animal Cognition*, *8*(4), 215–219. https://doi.org/10/bb8tsx

Over, H., & Carpenter, M. (2012). Putting the social into social learning: Explaining both selectivity and fidelity in children's copying behavior. *Journal of Comparative Psychology*, *126*(2), 182–192. https://doi.org/10/dg9n8q

Palmquist, C. M., Burns, H. E., & Jaswal, V. K. (2012). Pointing disrupts preschoolers' ability to discriminate between knowledgeable and ignorant informants. *Cognitive Development*, *27*(1), 54–63. https://doi.org/10/bxmdgn

Peirce, J. W., Gray, J. R., Simpson, S., Macaskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*, 195–203. https://doi.org/10/gft89w

Perner, J. (1991). *Understanding the representational mind*. MIT Press.

Perner, J., & Ruffman, T. (2005). Infants' insight into the mind: How deep? *Science*, *308*(5719), 214–216. https://doi.org/10/dcphjk

Pillow, B. H. (1989). Early understanding of perception as a source of knowledge. *Journal of Experimental Child Psychology*, *47*(1), 116–129. https://doi.org/10/czg4s7

Pillow, B. H., & Weed, S. T. (1997). Preschool children's use of information about age and perceptual access to infer another person's knowledge. *Journal of Genetic Psychology*, *158*(3), 365–376. https://doi.org/10/b72ngg

Pipek, P., Petrusková, T., Petrusek, A., Diblíková, L., Eaton, M. A., & Pyšek, P. (2018). Dialects of an invasive songbird are preserved in its invaded but not native source range. *Ecography*, *41*(2), 245–254. https://doi.org/10/gcxwtx

Povinelli, D., & DeBlois, S. (1992). Young children's (*Homo sapiens*) understanding of knowledge formation in themselves and others. *Journal of Comparative Psychology*, *106*(3), 228–238. https://doi.org/10/bz4jj6

Povinelli, D., Nelson, K. E., & Boysen, S. T. (1990). Inferences about guessing and knowing by chimpanzees (*Pan troglodytes*). *Journal of Comparative Psychology*, *104*(3), 203–210. https://doi.org/10/d53m2p

Povinelli, D., Parks, K. A., & Novak, M. A. (1991). Do rhesus monkeys (*Macaca mulatta*) attribute knowledge and ignorance to others? *Journal of Comparative Psychology*, *105*(4), 318–325. https://doi.org/10/dr8bf7

Povinelli, D., & Vonk, J. (2003). Chimpanzee minds: Suspiciously human? *Trends in Cognitive Sciences*, *7*(4), 157–160. https://doi.org/10/bzpgh5

Povinelli, D., & Vonk, J. (2004). We don't need a microscope to explore the chimpanzee's mind. *Mind & Language*, *19*(1), 1–28. https://doi.org/10/cz2mbc

Pratt, C., & Bryant, P. (1990). Young children understand that looking leads to knowing (so long as they are looking into a single barrel). *Child Development*, *61*(4), 973–982. https://doi.org/10/bt9n9h

Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, *4*, 515–526. https://doi.org/10/ddvt4n

Price, E. E., Wood, L. A., & Whiten, A. (2017). Adaptive cultural transmission biases in children and nonhuman primates. *Infant Behavior and Development*, *48*, 45–53. https://doi.org/10/djp4

Priewasser, B., Rafetseder, E., Gargitter, C., & Perner, J. (2018). Helping as an early indicator of a theory of mind: Mentalism or teleology? *Cognitive Development*, *46*, 69–78. https://doi.org/10/gd5n4d

Priewasser, B., Roessler, J., & Perner, J. (2013). Competition as rational action: Why young children cannot appreciate competitive games. *Journal of Experimental Child Psychology*, *116*(2), 545–559. https://doi.org/10/f3sktj

R Core Team. (2020). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. https://www.r-project.org/

Rakoczy, H., Hamann, K., Warneken, F., & Tomasello, M. (2010). Bigger knows better: Young children selectively learn rule games from adults rather than from peers. *British Journal of Developmental Psychology*, *28*(4), 785–798. https://doi.org/10/d653z3

Rakoczy, H., Warneken, F., & Tomasello, M. (2007). “This way!”, “No! That way!”—3-year olds know that two people can have mutually incompatible desires. *Cognitive Development*, *22*(1), 47–68. https://doi.org/10/fg4ng7

Rawlings, B., Flynn, E. G., & Kendal, R. (2017). To copy or to innovate? The role of personality and social networks in children’s learning strategies. *Child Development Perspectives*, *11*(1), 39–44. https://doi.org/10/f9396s

Reindl, E., Gwilliams, A. L., Dean, L. G., Kendal, R. L., & Tennie, C. (2020). Skills and motivations underlying children’s cumulative cultural learning: Case not closed. *Palgrave Communications*, *6*(106). https://doi.org/10/fmfq

Reindl, E., & Tennie, C. (2018). Young children fail to generate an additive ratchet effect in an open-ended construction task. *PLoS ONE*, *13*(6), 1–22. https://doi.org/10/gdqsxk

Rendell, L., Fogarty, L., Hoppitt, W. J. E., Morgan, T. J. H., Webster, M. M., & Laland, K. N. (2011). Cognitive culture: Theoretical and empirical insights into social learning strategies. *Trends in Cognitive Sciences*, *15*(2), 68–76. https://doi.org/10/dx8z7z

Renner, E., Kean, D., Atkinson, M., & Caldwell, C. A. (2021). The use of individual, social, and animated cue information by capuchin monkeys and children in a touchscreen task. *Scientific Reports*, *11*(1043). https://doi.org/10/fv57

Renner, E., Patterson, E. M., & Subiaul, F. (2020). Specialization in the vicarious learning of novel arbitrary sequences in humans but not orangutans. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *375*, 20190442. https://doi.org/10/fbwd

Repacholi, B. M., & Gopnik, A. (1997). Early reasoning about desires: Evidence from 14- and 18-month-olds. *Developmental Psychology*, *33*(1), 12–21. https://doi.org/10/dv6nfk

Robinson, E. J., Butterfill, S. A., & Nurmsoo, E. (2011). Gaining knowledge via other minds: Children's flexible trust in others as sources of information. *British Journal of Developmental Psychology*, *29*(4), 961–980. https://doi.org/10/czwr3k

Robinson, E. J., Haigh, S. N., & Pendle, J. E. C. (2008). Children's working understanding of the knowledge gained from seeing and feeling. *Developmental Science*, *11*(2), 299–305. https://doi.org/10/crvnqb

Rogers, A. R. (1988). Does biology constrain culture? *American Anthropologist*, *90*(4), 819–831. https://doi.org/10/bkrgxb

Rohwer, M., Kloo, D., & Perner, J. (2012). Escape from metaignorance: How children develop an understanding of their own lack of knowledge. *Child Development*, *83*(6), 1869–1883. https://doi.org/10/f3ssk9

Rostad, K., & Pexman, P. M. (2014). Developing appreciation for ambivalence: The understanding of concurrent conflicting desires in 4- to 7-year-old children. *Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale*, *68*(2), 122–132. https://doi.org/10/f6zz8g

Ruffman, T. K., Aitken, J., Wilson, A., Puri, A., & Taumoepeau, M. (2017). A re-examination of the broccoli task: Implications for children's understanding of subjective desire. *Cognitive Development*, *August*. https://doi.org/10/gd5n4z

Ruffman, T. K., & Olson, D. R. (1989). Children's ascriptions of knowledge to others. *Developmental Psychology*, *25*(4), 601–606. https://doi.org/10/bjfkh7

Saling, L. L., & Phillips, J. G. (2007). Automatic behaviour: Efficient not mindless. *Brain Research Bulletin*, *73*(1–3), 1–20. https://doi.org/10/dnw2g6

Sanz, C., Call, J., & Morgan, D. (2009). Design complexity in termite-fishing tools of chimpanzees (*Pan troglodytes*). *Biology Letters*, *5*(3), 293–296. https://doi.org/10/b8d923

Sasaki, T., & Biro, D. (2017). Cumulative culture can emerge from collective intelligence in animal groups. *Nature Communications*, *8*, 1–6. https://doi.org/10/f93j4k

Savage, R., Cornish, K., Manly, T., & Hollis, C. (2006). Cognitive processes in children's reading and attention: The role of working memory, divided attention, and response inhibition. *British Journal of Psychology*, *97*(3), 365–385. https://doi.org/10/fsp8cf

Saxton, M. (2017). *Child language: Acquisition and development*. Sage.

Schofield, D. P., McGrew, W. C., Takahashi, A., & Hirata, S. (2018). Cumulative culture in nonhumans: Overlooked findings from Japanese monkeys? *Primates*, *59*(2), 113–122. https://doi.org/10/gc7bdn

Senju, A., Southgate, V., Snape, C., Leonard, M., & Csibra, G. (2011). Do 18-month-olds really attribute mental states to others?: A critical test. *Psychological Science*, *22*(7), 878–880. https://doi.org/10/dbx7q8

Shea, N. (2013). Inherited representations are read in development. *British Journal for the Philosophy of Science*, *64*(1), 1–31. https://doi.org/10/fm6q

Shea, N., Boldt, A., Bang, D., Yeung, N., Heyes, C., & Frith, C. D. (2014). Supra-personal cognitive control and metacognition. *Trends in Cognitive Sciences*, *18*(4), 186–193. https://doi.org/10/f5xfn9

Smaldino, P. E., & Richerson, P. J. (2013). Human cumulative cultural evolution as a form of distributed computation. In P. Michelucci (Ed.), *Handbook of Human Computation* (pp. 979–992). Springer. https://doi.org/10/fmfr

So much more to know. (2005). *Science*, *309*(5731), 78–102. https://doi.org/10/c8qzkz

Sobel, D. M., & Kushnir, T. (2013). Knowledge matters: How children evaluate the reliability of testimony as a process of rational inference. *Psychological Review*, *120*(4), 779–797. https://doi.org/10/f5grwb

Sodian, B., Thoermer, C., & Dietrich, N. (2006). Two- to four-year-old children's differentiation of knowing and guessing in a non-verbal task. *European Journal of Developmental Psychology*, *3*(3), 222–237. https://doi.org/10/bt399d

Sodian, B., & Wimmer, H. (1987). Children's understanding of inference as a source of knowledge. *Child Development*, *58*(2), 424. https://doi.org/10/ctk9s6

Southgate, V., Senju, A., & Csibra, G. (2007). Action anticipation through attribution of false belief by 2-year-olds. *Psychological Science*, *18*(7), 587–592. https://doi.org/10/dz3qnx

Spaulding, S. (2020). What is mindreading? *Wiley Interdisciplinary Reviews: Cognitive Science*, *11*(3), 1–13. https://doi.org/10/gg8jhd

Stout, D. (2011). Stone toolmaking and the evolution of human culture and cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *366*(1567), 1050–1059. https://doi.org/10/c5ctq7

Subiaul, F., Zimmermann, L., Renner, E., Schilder, B., & Barr, R. (2016). Defining elemental imitation mechanisms: A comparison of cognitive and motor-spatial imitation learning across object- and computer-based tasks. *Journal of Cognition* *and Development*, *17*(2), 221–243. https://doi.org/10/d6jf

Tennie, C., Call, J., & Tomasello, M. (2009). Ratcheting up the ratchet: On the evolution of cumulative culture. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1528), 2405–2415. https://doi.org/10/fk2rxt

Tennie, C., Walter, V., Gampe, A., Carpenter, M., & Tomasello, M. (2014). Limitations to the cultural ratchet effect in young children. *Journal of Experimental Child Psychology*, *126*, 152–160. https://doi.org/10/f6c3bm

Terrier, N., Bernard, S., Mercier, H., & Clément, F. (2016). Visual access trumps gender in 3- and 4-year-old children's endorsement of testimony. *Journal of Experimental Child Psychology*, *146*, 223–230. https://doi.org/10/dr7p

Teufel, C., Alexis, D. M., Todd, H., Lawrance-Owen, A. J., Clayton, N. S., & Davis, G. (2009). Social cognition modulates the sensory coding of observed gaze direction. *Current Biology*, *19*(15), 1274–1277. https://doi.org/10/dq72d8

Teufel, C., Clayton, N. S., & Russell, J. W. (2013). Two-year-old children's understanding of visual perception and knowledge formation in others. *Journal of Cognition and Development*, *14*(2), 203–228. https://doi.org/10/djp5

Tomasello, M., Kruger, A. C., & Ratner, H. H. (1993). Cultural learning. *Behavioral and Brain Sciences*, *16*(03), 495. https://doi.org/10/fk5z77

Tong, Y., Wang, F., & Danovitch, J. (2020). The role of epistemic and social characteristics in children's selective trust: Three meta-analyses. *Developmental Science*, *23*(2), 1–18. https://doi.org/10/ggj3rh

UNESCO Institute for Statistics. (2017). *Literacy rates continue to rise from one generation to the next*. http://uis.unesco.org/sites/default/files/documents/fs45-literacy-rates-continue-rise-generation-to-next-en-2017_0.pdf

van Leeuwen, E. J. C., Cohen, E., Collier-Baker, E., Rapold, C., Schäfer, M., Schütte, S., & Haun, D. B. M. (2018). The development of human social learning across seven societies. *Nature Communications*, *9*(2076). https://doi.org/10/gdn9ch

VanderBorght, M., & Jaswal, V. K. (2009). Who knows best? Preschoolers sometimes prefer child informants over adult informants. *Infant and Child Development*, *18*, 61–71. https://doi.org/10.1002/icd

Vélez, N., & Gweon, H. (2019). Integrating incomplete information with imperfect advice. *Topics in Cognitive Science*, *11*(2), 1–17. https://doi.org/10/gfrdkv

Voelkl, B., & Huber, L. (2000). True imitation in marmosets. *Animal Behaviour*, *60*(2), 195–202. https://doi.org/10/fvftwr

Want, S. C., & Harris, P. L. (2001). Learning from other people's mistakes: Causal understanding in learning to use a tool. *Child Development*, *72*(2), 431–443. https://doi.org/10/cbvqfv

Warneken, F., & Tomasello, M. (2006). Altruistic helping in human infants and young chimpanzees. *Science*, *311*(5765), 1301–1303. https://doi.org/10/bmxtdd

Wasielewski, H. (2014). Imitation is necessary for cumulative cultural evolution in an unfamiliar, opaque task. *Human Nature*, *25*(1), 161–179. https://doi.org/10/gghf8x

Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory of mind development: The truth about false belief. *Child Development*, *72*(3), 655–684. https://doi.org/10/css9pc

Wellman, H. M., & Woolley, J. D. (1990). From simple desires to ordinary beliefs: The early development of everyday psychology. *Cognition*, *35*(3), 245–275. https://doi.org/10/c3xn46

Whalen, A., Griffiths, T. L., & Buchsbaum, D. (2018). Sensitivity to shared information in social learning. *Cognitive Science*, *42*(1), 168–187. https://doi.org/10/gcwvn7

Whiten, A. (2017a). Culture extends the scope of evolutionary biology in the great apes. *Proceedings of the National Academy of Sciences*, *114*(30), 7790–7797. https://doi.org/10/gbq9j2

Whiten, A. (2017b). Social learning and culture in child and chimpanzee. *Annual Review of Psychology*, *68*(1), 129–154. https://doi.org/10/djp7

Whiten, A., Custance, D. M., Gomez, J.-C., Teixidor, P., & Bard, K. A. (1996).

Imitative learning of artificial fruit processing in children (*Homo sapiens*) and chimpanzees (*Pan troglodytes*). *Journal of Comparative Psychology*, *110*(1), 3–14. https://doi.org/10/fvbhkp

Wilks, C. E. H., & Blakey, K. H. (2018). In the jungle of cultural complexity. *Evolutionary Anthropology*, *August*, 180–183. https://doi.org/10/d7hn

Wilks, C. E. H., Rafetseder, E., Renner, E., Atkinson, M., & Caldwell, C. A. (2021). Cognitive prerequisites for cumulative culture are context-dependent: Children's potential for ratcheting depends on cue longevity. *Journal of Experimental Child Psychology*, 105031. https://doi.org/10/fw5g

Wimmer, H., Hogrefe, G.-J., & Perner, J. (1988). Children's understanding of informational access as source of knowledge. *Child Development*, *59*(2), 386–396. https://doi.org/10/c649tg

Wimmer, H., & Perner, J. (1983). Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, *13*(1), 103–128. https://doi.org/10/cbrpxb

Wood, L. A., Harrison, R. A., Lucas, A. J., McGuigan, N., Burdett, E. R. R., & Whiten, A. (2016). "Model age-based" and "copy when uncertain" biases in children's social learning of a novel task. *Journal of Experimental Child Psychology*, *150*, 272–284. https://doi.org/10/f839t2

Wood, L. A., Kendal, R., & Flynn, E. G. (2012). Context-dependent model-based biases in cultural transmission: Children's imitation is affected by model age over model knowledge state. *Evolution and Human Behavior*, *33*(4), 387–394. https://doi.org/10/f337f6

Wood, L. A., Kendal, R., & Flynn, E. G. (2013a). Copy me or copy you? The effect of prior experience on social learning. *Cognition*, *127*(2), 203–213. https://doi.org/10/f4vh4v

Wood, L. A., Kendal, R., & Flynn, E. G. (2013b). Whom do children copy? Model-based biases in social learning. *Developmental Review*, *33*(4), 341–356. https://doi.org/10/f5hzms

Xu, F., Han, Y., Sabbagh, M. A., Wang, T., Ren, X., & Li, C. (2013). Developmental differences in the structure of executive function in middle childhood and adolescence. *PloS One*, *8*(10). https://doi.org/10/fmfs

Zelazo, P. D. (2006). The Dimensional Change Card Sort (DCCS): A method of assessing executive function in children. *Nature Protocols*, *1*(1), 297–301. https://doi.org/10/b9853h

Zelazo, P. D. (2015). Executive function: Reflection, iterative reprocessing, complexity, and the developing brain. *Developmental Review*, *38*, 55–68. https://doi.org/10/f7426r

Zelazo, P. D., Frye, D., & Rapus, T. (1996). An age-related dissociation between knowing rules and using them. *Cognitive Development*, *11*(1), 37–63. https://doi.org/10/c4m4r8

Zmyj, N., Buttelmann, D., Carpenter, M., & Daum, M. M. (2010). The reliability of a model influences 14-month-olds' imitation. *Journal of Experimental Child Psychology*, *106*(4), 208–220. https://doi.org/10/dw6t37

Zmyj, N., & Seehagen, S. (2013). The role of a model's age for young children's imitation: A research review. *Infant and Child Development*, *22*, 622–641. https://doi.org/10/djp8

Zwirner, E., & Thornton, A. (2015). Cognitive requirements of cumulative culture: Teaching is useful but not essential. *Scientific Reports*, *5*, 1–8. https://doi.org/10/f722rt

Appendices

Appendix A:
Verbal Script for Chapter 2

Same Goals Condition

1. "You are going to play a game with Rabbit."
2. "In the game you are going to be a chef, would you like to wear the chef hat?"
   E1 presents 1st set of buckets with carrots inside.
3. "Look, we have 2 buckets of carrots, in one of the buckets the carrots are orange inside, like this one (E1 shows orange carrot), and in the other bucket the carrots have worms inside, like this one (E1 shows worm carrot)."
4. "You are a chef, so you are looking for the carrots that have orange inside to put in your pan (E1 gives pan). And Rabbit likes to eat carrots that are orange inside, so Rabbit is looking for carrots with orange inside to put in the basket (E1 presents basket)."
5. "What is inside the carrots that you are looking for?"
   "What is inside the carrots that you are **not** looking for?"
6. "What is inside the carrots that Rabbit is looking for?"
   "What is inside the carrots that Rabbit is **not** looking for?"
7. "You and Rabbit both have 3 tokens; you need to pay a token into the box before you can pick a carrot from one of the buckets."
8. "It is Rabbit's turn first, Rabbit is going to pay a token into the box, and Rabbit is going to choose a bucket to pick a carrot from. Rabbit is going to have a peek inside the carrot and decide to keep it in the basket, or not keep it and give it to me."
9. "Now it is your turn."
   E1 guides child through their turn.
10. "You can pay a token into the box, and you can pick a carrot from one of the buckets, now you can have a peek inside and decide if you want to keep it and put it in your pan, or not keep it and give it to me."

Child is allowed to accept or reject either kind of carrot, E1 does not provide feedback on decisions.

11. "What was inside the carrot that Rabbit picked? Which bucket did they look in?"
12. "What was inside the carrot that you picked? Which bucket did you look in?"

    Child keeps any carrots they find in their pan; Rabbit's carrots stay in basket.

    E1 takes buckets out of sight, refills them, changes lid colours, and randomly swaps sides.
13. "Look you have 2 tokens left; would you like to play the game again?"

    Trial 2 follows same pattern as trial 1 (from step 4). Trial 3 follows trial 2 in same pattern. Game finishes after 3 trials.
14. "How were you deciding which buckets to choose?"
15. "Were you looking for the same kind of carrots as Rabbit, or a different kind of carrots?"

Different Goals Condition

1. "You are going to play a game with Bird."
2. "In the game you are going to be a chef, would you like to wear the chef hat?"

   E1 presents 1st set of buckets with carrots inside.
3. "Look, we have 2 buckets of carrots, in one of the buckets the carrots are orange inside, like this one (E1 shows orange carrot), and in the other bucket the carrots have worms inside, like this one (E1 shows worm carrot)."
4. "You are a chef, so you are looking for the carrots that have orange inside to put in your pan (E1 gives pan). And Bird likes to eat worms, so Bird is looking for carrots with worms inside to put in the basket (E1 presents basket)."
5. "What is inside the carrots that you are looking for?"

   "What is inside the carrots that you are **not** looking for?"
6. "What is inside the carrots that Bird is looking for?"

   "What is inside the carrots that Bird is **not** looking for?"
7. "You and Bird both have 3 tokens; you need to pay a token into the box before you can pick a carrot from one of the buckets."
8. "It is Bird's turn first, Bird is going to pay a token into the box, and Bird is going to choose a bucket to pick a carrot from. Bird is going to have a peek inside the carrot and decide to keep it in the basket, or not keep it and give it to me."

9. “Now it is your turn.”

   E1 guides child through their turn.

10. “You can pay a token into the box, and you can pick a carrot from one of the buckets, now you can have a peek inside and decide if you want to keep it and put it in your pan, or not keep it and give it to me.”

    Child is allowed to accept or reject either kind of carrot, E1 does not provide feedback on decisions.

11. “What was inside the carrot that Bird picked? Which bucket did they look in?”

12. “What was inside the carrot that you picked? Which bucket did you look in?”

    Child keeps any carrots they find in their pan; Bird’s carrots stay in basket.

    E1 takes buckets out of sight, refills them, changes lid colours, and randomly swaps sides.

13. “Look you have 2 tokens left; would you like to play the game again?”

    Trial 2 follows same pattern as trial 1 (from step 4). Trial 3 follows trial 2 in same pattern. Game finishes after 3 trials.

14. “How were you deciding which buckets to choose?”

15. “Were you looking for the same kind of carrots as Bird, or a different kind of carrots?”

Appendix B:
Verbal Script for Chapter 3

"Would you like to play a game watching some videos and trying to unlock some boxes?"

"Behind these pictures there are 4 locked boxes."

"When you touch the screen you will see one of these pictures, the picture tells you which box you are going to try to open."

"You can touch the screen to start now."

"Which box are you going to try and open?" (Turn screen away)

"Look, this is the box you are going to try to open, and here are 2 keys"

"One of the keys will unlock the padlock so you can open the box and see what's inside, but the other key will not unlock the padlock and the box will stay closed."

"Before you try to open your box, we are going to watch a video."

"You will see four pictures of other people who have boxes they want to unlock, you can choose one of the pictures to see a video of that person trying to open their box."

"You need to choose a video that will help you choose which key to try and open your box with."

(Turn screen back around)

"Touch the screen now to see the pictures, which video are you going to choose?"

After video: (Turn screen away)

"Now it's your turn to try and unlock your box. Which key are you going to choose?"

If they open the box:

"That's the right key. What's inside the box? You can put the duck in the pond!"

If they don't open the box:

"Oh no that's not the right key, the box won't open."

"Which key will open the box?"

"Would you like to try and unlock the other boxes?"

"You can touch the screen to start again." (Repeat for 3 more trials)

After all trials:

"Well done! You got _____ ducks in the pond."

"How were you deciding which videos to watch?"

"How were you deciding which keys to use?"

Appendix C:
Verbal Script for Chapter 4

"Look I have two types of frames. You can see through frames that are blue but you can't see through frames that are yellow." (Point to each type of frame when explaining)

"Look, can you see me through this blue frame? Can you see me through this yellow frame?" (Hold each frame in front of my face in turn)

"So which frame can you see me through? And which frame can I see you through?" (Record on tablet)

"Now it is time to play a finding and building game with Spots and Stripes."

"Look, one of these boxes has Lego inside, and one of the boxes is empty." (Show boxes)

"You and Spots and Stripes are going to try and find the Lego to build a tower on here." (Show Lego base)

"After I have hidden the Lego, I will ask you who you think will tell you where the Lego is."

"Before I hide the Lego, one frame will go in front of Spots, one frame will go in front of Stripes, and this frame will go in front of you." (Frames moved in front of demonstrators and then child)

"Now I will mix around the boxes and close the boxes so that the Lego is hidden."

"Now we can move the frames out of the way." (Child's frame removed from table, then demonstrator frames moved to the sides)

Pre-test trials: "Who do you think will tell you where the Lego is hidden?"

Test trials: "Who do you want to tell you where they think the Lego is hidden?" (Record response on tablet)

Pre-test trials: The informed demonstrator (blue frame) will then select the target box and reveal the Lego inside – ignoring child's selection.

Test trials: The demonstrator that the child selected will select the box – informed always picking Lego box, and uninformed always picking empty box.

Pre-test trials: "There's the Lego, you can start to build a tower on here."

Test trials (informed): "There's the Lego, you can start to build a tower on here." (Point at Lego base)

Test trials (uninformed): "Oh no, the Lego is not there."

All trials: "Let's have another go."

Start of test trials: "You are really good at guessing who can show you where the Lego is!"

"So this time you **choose who you want to tell you** where the Lego is? Try hard to pick the right person, or you might not get the Lego this time!"

Demonstrators will stay or switch sides randomly according to the tablet instructions. Repeated for 10 pre-test trials or until children are successful (select the informed demonstrator) in 5 consecutive trials.

"Look, one of these boxes has Lego inside, and one of the boxes is empty." (Show boxes)

Test trials: 1st: frames switch sides

2nd & 3rd: frame sides randomised

4th & 5th: no frames – informed faces table, uninformed faces away

After all trials are completed:

"Well done! You managed to find lots of Lego and you built a tower!"

"How were you choosing who you wanted to show you where the Lego was?"

Appendix D:
Pilot Study for Chapter 4

To evaluate the suitability of the task methodology, a pilot study was conducted with two- to seven-year-old children (*N* = 41) recruited at the University of Stirling Summer Science Festival. A number of variations of the intended task were tested to identify the most appropriate informants and age range, and to check the saliency of the available associative cues. Due to the need for participants to form an association, or infer a rule, prior to testing their capacity to use reasoned selective processes, it was necessary for each participant to experience a high number of pre-switch trials. However, this posed a problem due to the young age of the intended participants and the potential length of time they would need to remain focused and motivated to continue with the task. There were also queries surrounding whether the informants should be human adults or puppets, though it was expected that it would be easier for participants to ascribe perceptual access to human adults. To explore these methodological aspects of the task, we piloted versions of the task that employed puppets or adults as informants and varied the maximum number of pre-switch trials. As a result of the pilot testing, it was decided that the task should use human adults as informants which reduced the length of each trial and appeared to hold participants' attention longer. By employing adult informants, the number of pre-switch trials could be set at the higher number to provide more opportunity to form an association, or infer a rule, prior to the switch. In the pilot study, coloured wooden shapes were used as tokens to indicate which box each informant thought the Lego was inside (similarly to non-verbal theory of mind tasks; e.g., Call & Tomasello, 1999); children were then asked which box they wanted to open to try and find the Lego. We decided in the main experiment to modify this aspect of the task by removing the tokens and, rather than asking children to choose a box, asking them to choose an informant as this appears to be a more suitable measure to ascertain whether children recognise who has access to knowledge.